



Red Hat Satellite 6.2

Server Administration Guide

Administering a Red Hat Satellite 6 Server.

Edition 1.0

Last Updated: 2018-07-25

Red Hat Satellite 6.2 Server Administration Guide

Administering a Red Hat Satellite 6 Server.
Edition 1.0

Red Hat Satellite Documentation Team
satellite-doc-list@redhat.com

Legal Notice

Copyright © 2016 Red Hat.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The Red Hat Satellite 6 Server Administration Guide provides instructions on how to configure and administer a Red Hat Satellite 6 Server. Before continuing with this workflow you must have successfully installed a Red Hat Satellite 6 Server and any required Capsule Servers.

Table of Contents

CHAPTER 1. ACCESSING RED HAT SATELLITE	3
1.1. LOGGING IN TO RED HAT SATELLITE	3
1.2. CHANGING THE PASSWORD IN RED HAT SATELLITE	5
CHAPTER 2. STARTING AND STOPPING RED HAT SATELLITE	7
CHAPTER 3. CONFIGURING ORGANIZATIONS, LOCATIONS AND LIFE CYCLE ENVIRONMENTS	8
3.1. ORGANIZATIONS	8
3.2. LOCATIONS	12
3.3. LIFE CYCLE ENVIRONMENTS	13
3.4. VIEWING IMPORT HISTORY	17
CHAPTER 4. USERS AND ROLES	18
4.1. CREATING AND MANAGING USERS	18
4.2. CREATING USER GROUPS	21
4.3. CREATING AND MANAGING ROLES	22
4.4. GRANULAR PERMISSION FILTERING	25
CHAPTER 5. BACKUP AND DISASTER RECOVERY	28
5.1. BACKING UP SATELLITE SERVER OR CAPSULE SERVER	28
5.2. RESTORING SATELLITE SERVER OR CAPSULE SERVER FROM A BACKUP	33
5.3. BACKING UP AND RESTORING CAPSULE SERVER USING A SNAPSHOT	35
5.4. RENAMING A SATELLITE SERVER OR CAPSULE SERVER	35
CHAPTER 6. MAINTAINING A RED HAT SATELLITE SERVER	40
6.1. LOGGING AND REPORTING	40
6.2. ENABLING DEBUG LOGGING	41
6.3. COLLECTING INFORMATION FROM LOG FILES	42
6.4. USING LOG FILES IN SUPPORT CASES	43
6.5. ACCESSING CUSTOMER PORTAL SERVICES FROM RED HAT SATELLITE	43
6.6. USING RED HAT INSIGHTS WITH SATELLITE SERVER	47
6.7. MONITORING SATELLITE SERVER IN THE WEB UI	47
CHAPTER 7. MONITORING CAPSULE SERVERS	48
7.1. VIEWING GENERAL CAPSULE INFORMATION	48
7.2. MONITORING SERVICES	48
7.3. MONITORING PUPPET	48
CHAPTER 8. CONFIGURING EXTERNAL AUTHENTICATION	50
8.1. USING LDAP	50
8.2. USING IDENTITY MANAGEMENT	55
8.3. USING ACTIVE DIRECTORY	57
8.4. CONFIGURING EXTERNAL USER GROUPS	61
8.5. EXTERNAL AUTHENTICATION FOR PROVISIONED HOSTS	62
CHAPTER 9. CUSTOMIZING SATELLITE SERVER	67
9.1. ADDING ADDITIONAL PLUG-INS	67
9.2. USING FOREMAN HOOKS	68

CHAPTER 1. ACCESSING RED HAT SATELLITE

1.1. LOGGING IN TO RED HAT SATELLITE

After Red Hat Satellite has been installed and configured, use the web user interface to log in to Satellite for further configuration.

Procedure 1.1. Installing the Katello Root CA Certificate

The first time that you log in to Satellite, it is possible that you will see a warning informing you that you are using the default self-signed certificate. It is possible that you will not be able to connect this browser to Satellite until the proper root CA certificate is installed in the browser. Use the following procedure to locate the root CA certificate on the Satellite server and to install it in your browser.

1. Browse to **`http://HOSTNAME/pub`**
2. Select *katello-server-ca.crt*.
3. Import the certificate into your browser.

Procedure 1.2. To Log in to Satellite:

1. Access the Satellite Server using a web browser pointed to the following address:

`https://HOSTNAME/`

To identify your host name, use the **hostname** command at the prompt:

```
# hostname
```



IMPORTANT

An untrusted connection warning appears on your web browser when accessing Satellite for the first time. Accept the self-signed certificate and add the Satellite URL as a security exception to override the settings. This procedure might differ depending on the browser being used.

Only do this if you are sure that the Satellite URL is a trusted source.




Figure 1.1. Untrusted Connection Warning

2. Enter the user name and password created during the configuration process. If a user was not created during the configuration process, the default user name is *admin*. If you have forgotten the password of the default administrative account, *admin*, see [Section 1.2, “Changing the Password in Red Hat Satellite”](#).

Result

When you have successfully logged in, you are taken to the Satellite dashboard. The dashboard contains an overview of the Satellite and the hosts registered. For more information, see [Using the Red Hat Satellite Content Dashboard](#) and [Searching and Bookmarking](#) in the *Red Hat Host Configuration Guide*.

The main navigation tabs are as follows:

Table 1.1. Navigation Tabs

Navigation Tabs	Description
Default Organization	Clicking this tab changes the organization and location. If no organization or location is selected, the default organization is <i>Any Organization</i> and the default location is <i>Any Location</i> . Use this tab to change to different values.
Monitor	Provides summary dashboards and reports.

Navigation Tabs	Description
Content	Provides content management tools. This includes Content Views, Activation Keys, and Life Cycle Environments.
Containers	Provides container management tools.
Hosts	Provides host inventory and provisioning configuration tools.
Configure	Provides general configuration tools and data including Host Groups and Puppet data.
Infrastructure	Provides tools on configuring how Satellite 6 interacts with the environment.
Red Hat Insights	Provides Red Hat Insights management tools.
Administer	Provides advanced configuration for settings such as Users and RBAC, as well as general settings.
User Name	Provides user administration where users can edit their personal information.



WARNING

Satellite Server will be listed as a host itself even if it is not self-registered. Do not delete the Satellite Server from the list of hosts.

1.2. CHANGING THE PASSWORD IN RED HAT SATELLITE

These steps show how to change your password.

Procedure 1.3. To Change Your Red Hat Satellite Password:

1. Click your user name at the top right corner.
2. Select **My Account** from the menu.
3. In the **Password** field, type in a new password.
4. In the **Verify** field, type in the new password again.
5. Click the **Submit** button to save your new password.

Resetting the Password in Red Hat Satellite

If you have forgotten the administrative password, log on to the Satellite command-line interface to reset it:

```
# foreman-rake permissions:reset  
Reset to user: admin, password: qwJxBptxb7Gfcjj5
```

This will reset the password of the default user *admin* to the one printed on the command line. Change this password upon logging in to prevent any security issues from occurring.

CHAPTER 2. STARTING AND STOPPING RED HAT SATELLITE

Satellite provides the **katello-service** command to manage Satellite services from the command line. This is useful when upgrading Satellite or when creating a backup, see the [Red Hat Satellite Installation Guide](#) for details on these use cases.

After installing Satellite with the **satellite-installer** command, all Satellite services are started and enabled automatically. View the list of these services by executing:

```
# katello-service list
```

To see the status of running services, execute:

```
# katello-service status
```

To stop all Satellite services, execute:

```
# katello-service stop
```

To start all Satellite services, execute:

```
# katello-service start
```

To restart all Satellite services, execute:

```
# katello-service restart
```

CHAPTER 3. CONFIGURING ORGANIZATIONS, LOCATIONS AND LIFE CYCLE ENVIRONMENTS

Red Hat Satellite 6 takes a consolidated approach to Organization and Location management. System administrators define multiple Organizations and multiple Locations in a single Satellite Server. For example, a company might have three Organizations (Finance, Marketing, and Sales) across three countries (United States, United Kingdom, and Japan). In this example, the Satellite Server manages all Organizations across all geographical Locations, creating nine distinct contexts for managing systems. In addition, users can define specific locations and nest them to create a hierarchy. For example, Satellite administrators might divide the United States into specific cities, such as Boston, Phoenix, or San Francisco.

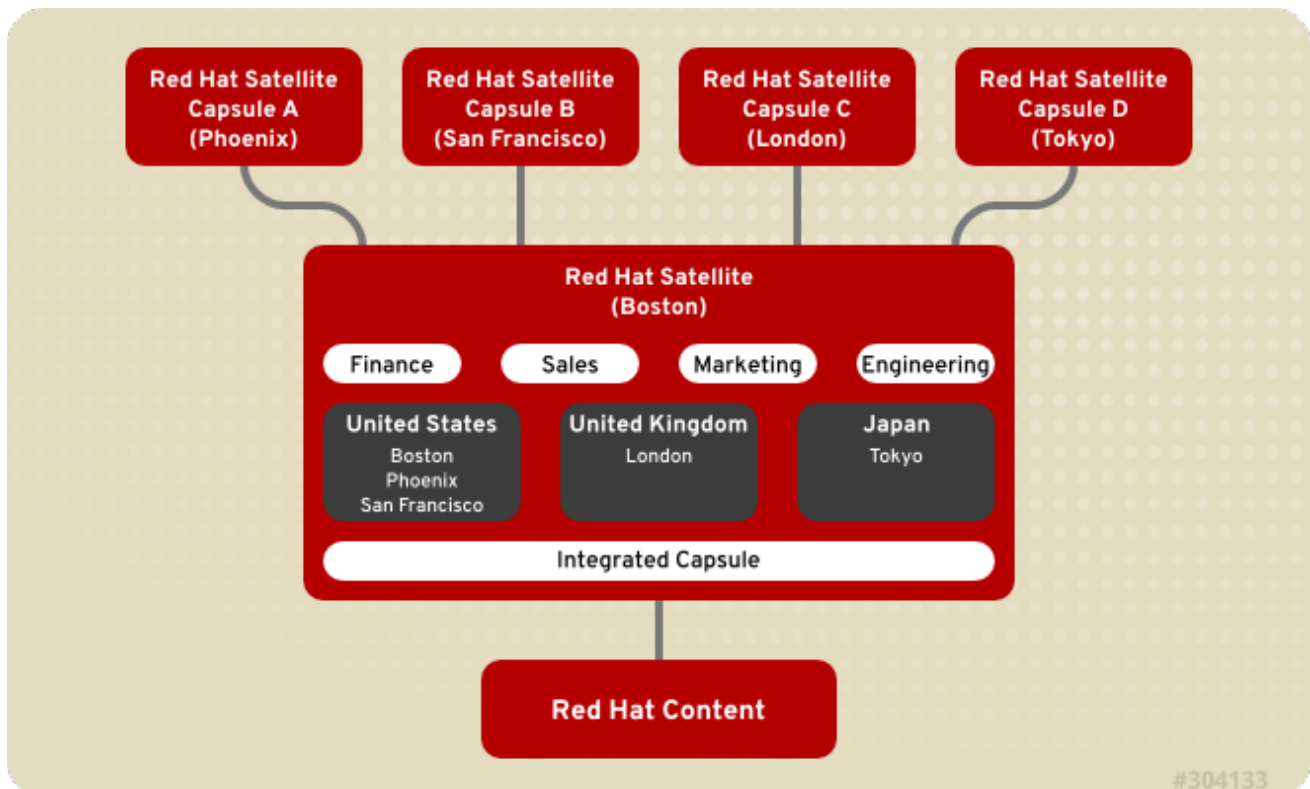


Figure 3.1. Example Topology for Red Hat Satellite 6

The main Satellite Server retains the management function, while the content and configuration is synchronized between the main Satellite Server and a Satellite Capsule Server assigned to certain locations.


3.1. ORGANIZATIONS

Organizations divide hosts into logical groups based on ownership, purpose, content, security level, or other divisions.

Multiple organizations can be viewed, created, and managed within the web UI. Software and host entitlements can be allocated across many organizations, and access to those organizations controlled.

Each organization must be created and used by a single Red Hat customer account, however each account can manage multiple organizations. Subscription Manifests can only be imported into a single organization and Satellite will not upload a certificate that has already been uploaded into a different organization.

The Red Hat Satellite installation process creates an organization called **Default Organization** unless another name is specified. The organization name has a corresponding label.



IMPORTANT

If a new user is not assigned a default organization their access will be limited. To grant systems rights to users, assign them to a default organization and have them log out and log back in again.

3.1.1. Creating an Organization

These steps show how to create a new organization.

Procedure 3.1. To Create an Organization:

1. Navigate to **Administer** → **Organizations**.
2. Click **New Organization**.
3. In the **Name** field, insert the name of the new organization.
4. In the **Label** field, insert the label of the new organization.
5. In the **Description** field, insert a description of the new organization.
6. Click **Submit**.
7. Select the hosts to assign to the new organization.
 - Click **Assign All** to assign all hosts with no organization to the new organization.
 - Click **Manually Assign** to manually select and assign the hosts with no organization.
 - Click **Proceed to Edit** to skip assigning hosts.
8. Specify the configuration details of the organization such as Capsule Servers, subnets or compute resources. You can modify these settings later as described in [Section 3.1.4, “Editing an Organization”](#).
9. Click **Submit**.

3.1.2. Creating an Organization Debug Certificate

These steps show how to generate and download a debug certificate for an organization. Debug certificates enable you to browse all content from an organization's repositories and are required for exporting provisioning templates.

Procedure 3.2. To Create a New Organization Debug Certificate:

1. Navigate to **Administer** → **Organizations**.
2. Select an organization for which you want to generate a debug certificate.

3. Click **Generate and Download**. This generates a debug certificate.
4. Save the certificate file in a secure location.

**NOTE**

Debug Certificates are automatically generated for provisioning template downloads if they do not already exist in the organization for which they are being downloaded.

3.1.3. Using an Organization Debug Certificate

You can view an organization's repository content using a browser or using the API if you have a debug certificate for that organization. The previous section describes creating and downloading the certificate which is in the X.509 format. To use a browser you must first convert the X.509 certificate to a format your browser supports and then import the certificate. The `curl` utility only requires extracting the certificate and key into separate files.

Procedure 3.3. To Use an Organization Debug Certificate in Firefox:

1. Create and download an organization certificate as described in [Procedure 3.2, "To Create a New Organization Debug Certificate:"](#).
2. Open the X.509 certificate, for example, for the default organization:

```
$ vi 'Default Organization-key-cert.pem'
```

3. Copy the contents of the file from `-----BEGIN RSA PRIVATE KEY-----` to `-----END RSA PRIVATE KEY-----` inclusive, into a file called `key.pem`.
4. Copy the contents of the file from `-----BEGIN CERTIFICATE-----` to `-----END CERTIFICATE-----` inclusive, into a file called `cert.pem`.
5. Enter a command as follows to create a PKCS12 format certificate and enter a password or phrase when prompted:

```
$ openssl pkcs12 -keypbe PBE-SHA1-3DES -certpbe PBE-SHA1-3DES -
export -in cert.pem -inkey key.pem -out organization_label.pfx -name
'organization_name'
Enter Export Password:
Verifying - Enter Export Password:
```

6. Using the preferences tab, import the resulting `pfx` file into your browser: Navigate to **Edit** → **Preferences** → **Advanced Tab**. Select **View Certificates** in the **Certificates** view to open the **Certificate Manager**. On the **Your Certificates** tab, click **Import** and select the `pfx` file to load. You will be prompted for the password or phrase used when making the certificate.
7. Enter a URL in the following format into your browser's address bar to begin browsing for repositories:

```
http://satellite.example.com/pulp/repos/organization_label
```

Pulp uses the organization label so the URL must use the organization label too.

Procedure 3.4. To Use an Organization Debug Certificate with curl:

1. Create and download an organization certificate as described in [Procedure 3.2, “To Create a New Organization Debug Certificate:”](#).
2. Open the X.509 certificate, for example, for the default organization:

```
$ vi 'Default Organization-key-cert.pem'
```

3. Copy the contents of the file from -----BEGIN RSA PRIVATE KEY----- to -----END RSA PRIVATE KEY----- inclusive, into a file called **key.pem**.
4. Copy the contents of the file from -----BEGIN CERTIFICATE----- to -----END CERTIFICATE----- inclusive, into a file called **cert.pem**.
5. Find a valid URL for a repository. You can use the browsing method described in the previous procedure or use the web UI. For example, using the web UI, navigate to **Content** → **Products** and select a Product by name. On the **Repositories** tab, select a repository by name and look for the **Published At** entry.
6. To use **curl** to access a repository, use a command as follows:

```
$ curl -k --cert cert.pem --key key.pem
http://satellite.example.com/pulp/repos/Default_Organization/Library
/content/dist/rhel/server/7/7Server/x86_64/sat-tools/6.2/os/
```

Ensure the paths to **cert.pem** and **key.pem** are the correct absolute paths otherwise the command will fail silently.

3.1.4. Editing an Organization

Procedure 3.5. To Edit an Organization:

1. Navigate to **Administer** → **Organizations**.
2. Click the name of the organization to be edited.
3. Select the resource to edit from the list on the left.
4. Click the name of the desired items to add them to the **Selected Items** list.
5. Click **Submit**.



NOTE


Users with administrator privileges are not listed under the **Users** tab when editing an organization.

3.1.5. Removing an Organization

Procedure 3.6. To Remove an Organization:

You can delete an organization if the organization is not associated with any life cycle environments or host groups. If there are any life cycle environments or host groups associated with the organization you are about to delete, deselect them by going to **Organizations** and clicking the relevant tabs. It is not recommended to delete the default organization created during installation because the default organization is a placeholder for any unassociated hosts in the Satellite environment. There must be at least one organization in the environment in any given time.

1. Navigate to **Administer** → **Organizations**.
2. Select **Delete** from the drop-down menu to the right of the name of the organization you want to remove.
3. An alert box appears:




Delete *Organization*?

4. Click **OK** to delete the organization.

3.2. LOCATIONS

Locations divide organizations into logical groups based on geographical location. Each location is created and used by a single Red Hat customer account, although each account can manage multiple locations and organizations.

The Red Hat Satellite installation process creates a location called **Default Location** unless another name is specified. If a new user is not assigned a default location their access will be limited. To grant system rights to users, assign a default location and have them log out and log in again.



IMPORTANT

You cannot delete the default location, but you can rename it to suit your needs. Satellite returns an error message if you try to delete the default location using either the web UI or the command line.

3.2.1. Creating a Location

These steps show how to create a location.

Procedure 3.7. To Create a Location:

1. Navigate to **Administer** → **Locations**.
2. Click **New Location**.
3. Insert the name of the new location in the **Name** field. If you want to create a nested location, select a **Parent** location from the drop-down menu. Optionally, specify a **Description** of the location. Click **Submit**.
4. Select the hosts to assign to the new location.
 - Click **Assign All** to assign all hosts with no location to the new location.
 - Click **Manually Assign** to manually select and assign the hosts with no location.

- Click **Proceed to Edit** to skip assigning hosts.
- 5. Specify the configuration details of the location such as Capsule Servers, subnets or compute resources. You can modify these settings later as described in [Section 3.2.2, “Editing a Location”](#).
- 6. Click **Submit**.

3.2.2. Editing a Location

Procedure 3.8. To Edit a Location:

1. Navigate to **Administer** → **Locations**.
2. Click the name of the location to be edited.
3. Select the resource to edit from the list on the left.
4. Click the name of the desired items to add them to the **Selected Items** list.
5. Click **Submit**.

3.2.3. Removing a Location

These steps show how to remove an existing location. Deleting the default location created during installation is currently not supported.

Procedure 3.9. To Remove a Location:

1. Navigate to **Administer** → **Locations**.
2. Select **Delete** from the drop-down menu to the right of the name of the location you want to remove.

An alert box appears:

An alert box dialog with a vertical bar on the left and the text "Delete Location?" inside.

3. Click **OK**.

3.3. LIFE CYCLE ENVIRONMENTS

Application life cycles are divided into *life cycle environments*, which represent each stage of the application life cycle. Life cycle environments are linked to form an *environment path*. You can promote content along the environment path to the next life cycle environment when required. For example, if development ends on a particular version of an application, you can promote this version to the testing environment and start development on the next version.

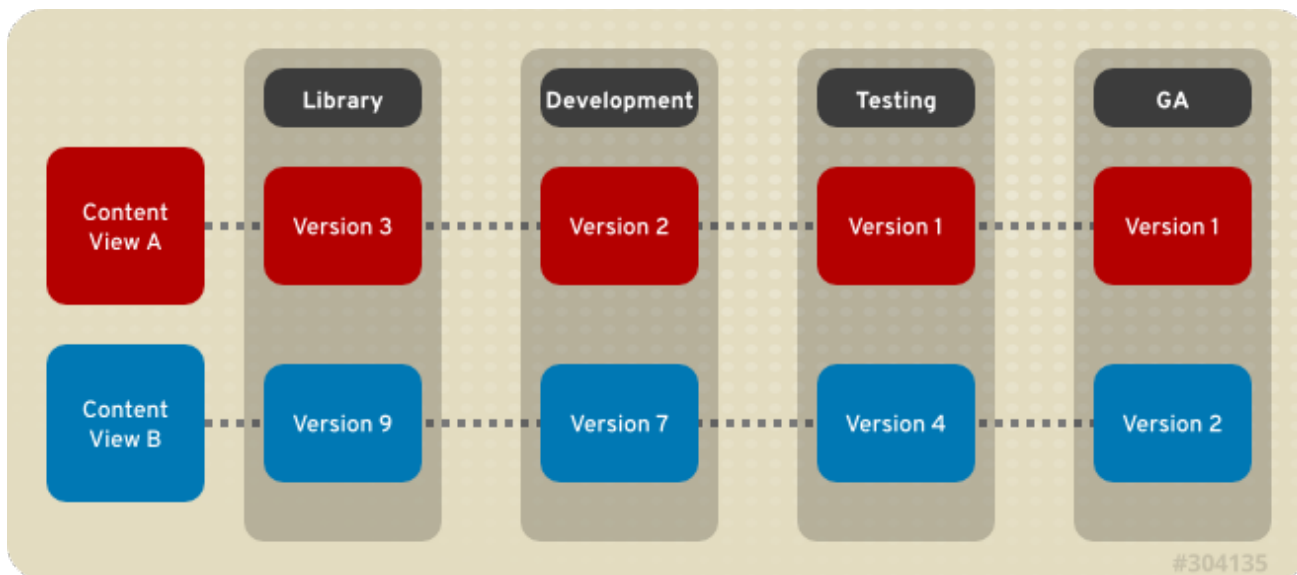


Figure 3.2. An Environment Path Containing Four Environments

3.3.1. Creating Life Cycle Environments

This procedure describes how to create a life cycle environment in Red Hat Satellite.

Procedure 3.10. To Create a Life Cycle Environment:

1. Select an organization from the menu in the top left hand corner.
2. Click **Content** → **Life Cycle Environments** and then click **New Environment Path**.
3. Insert a name and a label (automatically fills in the **Name** field input) for the life cycle environment. The **Description** field is optional.
4. Click **Save** to create the environment.

3.3.2. Adding Life Cycle Environments to a Red Hat Satellite Capsule Server

If a newly created Red Hat Satellite Capsule Server has content functionality enabled, the Capsule Server needs an environment added to enable it to synchronize content from the Satellite Server and provide content to host systems.



IMPORTANT

The Satellite Capsule Server is configured through the Satellite Server's command line interface (CLI). Execute all **hammer** commands on the Satellite Server.

Procedure 3.11. To Add Environments to the Satellite Capsule Server:

1. Log in to the Satellite Server CLI as root.
2. Choose the desired Red Hat Satellite Capsule Server from the list and take note of its **id**:

```
# hammer capsule list
```

The Satellite Capsule Server's details can be verified using the command:

```
# hammer capsule info --id capsule_id_number
```

3. Verify the list of life cycle environments available for the Red Hat Capsule Server and note down the **environment id**:

```
# hammer capsule content available-lifecycle-environments --id capsule_id_number
```

Where:

- **available-lifecycle-environments** are life cycle environments that are available to the Satellite Capsule but are currently not attached to the Satellite Capsule.

4. Add the life cycle environment to the Satellite Capsule Server:

```
# hammer capsule content add-lifecycle-environment --id capsule_id_number --environment-id environment_id_number
```

Where:

- *capsule_id_number* stands for the Satellite Capsule Server's identification number.
- *environment_id_number* stands for the life cycle environment's identification number.

Repeat this step for every life cycle environment to be added to the Capsule Server.

5. Synchronize the content from the Satellite Server's environment to the Satellite Capsule Server:

```
# hammer capsule content synchronize --id capsule_id_number
```

When an external Satellite Capsule Server has various life cycle environments, and only one life cycle environment needs to be synchronized, it is possible to target a specific environment by specifying the environment identification number:

```
# hammer capsule content synchronize --id external_capsule_id_number --environment-id environment_id_number
```

3.3.3. Promoting Content Views

After you have created a Content View and an environment path consisting of two or more life cycle environments, you can promote the Content View from one environment to the next as required. This means that the most recent version of the Content View that exists in a specified environment will be promoted, or copied, to the next environment in the life cycle environment path.

You can promote a Content View to any environment where that version does not exist. The system automatically suggests the next environment in the life cycle environment path, but you can override this and promote to a different environment if required.

Procedure 3.12. To Promote a Content View:

1. On the main menu, click **Content** → **Content Views**.
2. In the **Name** column, click the name of the Content View that you want to promote.
3. On the **Versions** tab, identify the latest version, and click **Promote**.
4. Identify the promotion path where you want to promote the Content View, select the appropriate life cycle environment, and click **Promote Version**.
5. After the promotion has completed, the **Versions** tab updates to display the new status of your Content Views.

3.3.4. Removing Life Cycle Environments From Satellite Server

This procedure describes how to remove a life cycle environment from Red Hat Satellite.

Procedure 3.13. To Remove a Life Cycle Environment:

1. On the main menu, click **Content** → **Life Cycle Environments**.
2. Click the name of the life cycle environment that you want to remove, and then click **Remove Environment**.
3. In the confirmation dialog box, click **Remove** to remove the environment.

**NOTE**

You can only delete the latest environment in an environment path. For example, if three environments exist in the order **Library**, **Dev**, and **Prod**, you need to delete **Prod** before you can delete **Dev**. You cannot delete the **Library** environment.

3.3.5. Removing Life Cycle Environments from Capsule Server

There are multiple reasons to remove life cycle environments from Capsule Server. For example:

- When life cycle environments are no longer relevant to the host systems
- When life cycle environments have been incorrectly added to Capsule Server

Procedure 3.14. To remove a life cycle environment from Capsule Server:

1. Log in to the Satellite Server CLI as the root user.
2. Choose the desired Capsule Server from the list and take note of its **id**:

```
# hammer capsule list
```

The Capsule Server's details can be verified using the command:

```
# hammer capsule info --id capsule_id_number
```

3. Verify the list of life cycle environments currently attached to the Capsule Server and take note of the **environment id**:

```
# hammer capsule content lifecycle-environments --id  
capsule_id_number
```

4. Remove the life cycle environment from Capsule Server:

```
# hammer capsule content remove-lifecycle-environment --id  
capsule_id_number --environment-id environment_id
```

Where:

- *capsule_id_number* is Capsule Server's identification number.
- *environment_id* is the life cycle environment's identification number.

Repeat this step for every life cycle environment to be removed from the Capsule Server.

5. Synchronize the content from the Satellite Server's environment to Capsule Server:

```
# hammer capsule content synchronize --id capsule_id_number
```

3.4. VIEWING IMPORT HISTORY

These steps show how to view an import history in Red Hat Satellite.

Procedure 3.15. To View an Import History:

1. Click **Content** → **Red Hat Subscriptions**.
2. Click the **Manage Manifest** button.
3. Click the **Import History** tab.

CHAPTER 4. USERS AND ROLES

A *User* defines a set of details for individuals using the system. Users can be associated with organizations and environments, so that when they create new entities, the default settings are automatically used. Users can also have one or more *roles* attached, which grants them rights to view and manage organizations and environments. See [Section 4.1, “Creating and Managing Users”](#) for more information on working with users.

You can manage permissions of several users at once by organizing them into *user groups*. User groups themselves can be further grouped to create a hierarchy of permissions. See [Section 4.2, “Creating User Groups”](#) for more information on creating user groups.

Roles define a set of permissions and access levels. Each role contains one or more *permission filters* that specify the *actions* allowed for the role. Actions are grouped according to the *Resource type*. Once a role has been created, users and user groups can be associated with that role. This way, you can assign the same set of permissions to large groups of users. Red Hat Satellite provides a set of predefined roles and also enables creating custom roles and permission filters as described in [Section 4.3, “Creating and Managing Roles”](#).

4.1. CREATING AND MANAGING USERS

For the administrator, Red Hat Satellite provides the ability to create, modify, and remove users. Also, it is possible to configure access permissions through assigning roles to users.

4.1.1. Creating a User

The following steps show how to create a user:

Procedure 4.1. To Create a User:

1. Navigate to **Administer** → **Users** and then click **New User**.
2. On the **User** tab, enter the required details.
3. On the **Locations** tab, select the required locations for this user.
4. On the **Organizations** tab, select organizations accessible to this user. The current active organization is selected by default. If you specify multiple organizations, you can select the default organization for user login from the drop-down list.
5. On the **Roles** tab, select the required roles for this user. Active roles are displayed in the right panel.
6. Click **Submit** to create the user.

4.1.2. Editing a User

The following steps show how to edit details of an existing user:

Procedure 4.2. To Edit an Existing User:

1. Navigate to **Administer** → **Users**.

2. Click the user name of the user to be altered. General information about the user will appear on the right.
3. In the **User** tab, you can modify the user's user name, first name, surname, email address, default location, default organization, language, and password.
4. In the **Locations** tab, you can modify the user's assigned locations.
5. In the **Organizations** tab, you can modify the user's assigned organizations.
6. In the **Roles** tab, you can modify the user's assigned roles.
7. Click **Save** to save your changes.

4.1.3. Assigning Roles to a User

By default, a new user has no roles assigned. The following procedure describes how to assign one or more roles to a user. You can select from predefined roles, or define a custom role as described in [Section 4.3.1, "Creating a Role"](#). You can apply a similar procedure to user groups.


Procedure 4.3. To Assign a Role to a User:

1. Navigate to **Administer** → **Users**. If a user account created is not listed, check that you are currently viewing the right organization. To list all users in Satellite, click **Default Organization** and then **Any Organization**. The organization view is changed to **Any Context**.
2. Click the user name of the user that you want to modify. General information about the user appears on the right.
3. Click the **Locations** tab, and select a location if none is assigned.
4. Click the **Organizations** tab, and check that an organization is assigned.
5. Click the **Roles** tab to display the list of available role assignments.
6. Select role you want to assign to the user in the **Roles** list. The list contains the predefined roles, as well as any custom roles, see [Table 4.1, "Predefined Roles Available in Red Hat Satellite"](#). Alternatively, select the **Administrator** check box to assign all available permissions to the selected user.
7. Click **Save**.

To view the roles assigned to any user, click the **Roles** tab; the assigned roles are listed under **Selected items**. To remove a role, from the **Selected items**, click a role name and it will be removed.

4.1.4. Configuring Email Notifications

Email notification is a per-user setting, with no email notifications enabled by default. If you want email notifications sent to a group's email address, instead of an individual's email address, create a user account with the group's email address and minimal Satellite permissions, then subscribe the user account to the desired notification types.



IMPORTANT

Satellite Server does not enable outgoing emails by default, therefore you must review your email configuration. For more information, see [Configuring Satellite Server for Outgoing Emails](#) in the *Red Hat Satellite Installation Guide*

Procedure 4.4. To Configure Email Notifications:

1. Navigate to **Administer** → **Users**.
2. Click the **Username** of the user you want to edit.
3. On the **User** tab, check the **Email address** field. Ensure that it contains a valid email address. The address will be associated with the user account, and the notifications selected in the following steps will be sent there.
4. Click the **Email Preferences** tab and select **Mail enabled** to enable email notifications.
5. Select the notifications you want the user to receive.
 - **Audit summary** is a summary of all activity audited by the Satellite Server. To enable these notifications, select the frequency of emails from the drop-down list that offers **Daily**, **Weekly**, or **Monthly** updates. Enter a query in the associated query field to narrow the audit activity included.
 - **Host built** is a notification sent when a host is built. To enable these notifications, select **Subscribe** from the drop-down menu.
 - **Host errata advisory** is a summary of applicable and installable errata for hosts managed by the user. To enable these notifications, select the frequency of emails from the drop-down list that offers **Daily**, **Weekly**, or **Monthly** updates.
 - **OpenSCAP policy summary** is a summary of OpenSCAP policy reports and their results. To enable these notifications, select the frequency of emails from the drop-down list that offers **Daily**, **Weekly**, or **Monthly** updates.
 - **Promote errata** is a notification sent only after a Content View promotion. It contains a summary of errata applicable and installable to hosts registered to the promoted Content View. This allows you to monitor what updates have been applied to which hosts. To enable these notifications, select **Subscribe** from the drop-down menu.
 - **Puppet error state** is a notification sent after a host reports an error related to Puppet. To enable these notifications, select **Subscribe** from the drop-down menu.
 - **Puppet summary** is a summary of Puppet reports. To enable these notifications, select the frequency of emails from the drop-down list that offers **Daily**, **Weekly**, or **Monthly** updates.
 - **Sync errata** is a notification sent only after synchronizing a repository. It contains a summary of new errata introduced by the synchronization. To enable these notifications, select **Subscribe** from the drop-down menu.
6. Click **Submit**.

Testing Email Delivery

To test email delivery to the email address associated with a user account, open the Satellite web UI, navigate to **Administer** → **Users**, click on the user name, click the **Email Preferences** tab and click **Test email**. A test email message is then sent immediately to the user's email address. If it does not arrive, first verify the user's email address, then the Satellite Server's email configuration, after which you may need to examine firewall and mail server logs.

Testing Email Notifications

To verify that your subscription to selected email notifications is valid, you can have all periodic notifications sent to you on request. Note that it will trigger all notifications scheduled for the specified frequency, and affect all users who have subscribed to it. Sending on request notifications to individual users is currently not supported.

To trigger the notifications, execute the following command on the Satellite Server:

```
# foreman-rake reports:frequency
```

Where *frequency* stands for a specific time period:

- daily
- weekly
- monthly

4.1.5. Removing a User

The following procedure describes how to remove an existing user.

Procedure 4.5. To Remove a User:

1. On the main menu, click **Administer** → **Users** to open the **Users** page.
2. Click the **Delete** link to the right of the user name you want to delete.
3. In the alert box, click **OK** to delete the user.

4.2. CREATING USER GROUPS

With Red Hat Satellite, you can assign permissions to groups of users. You can also create user groups as collections of other user groups. If using an external authentication source, you can map Satellite user groups to external user groups as described in [Section 8.4, “Configuring External User Groups”](#).

User groups are defined in an organizational context, meaning that you must select an organization before you can access user groups.

Procedure 4.6. To Create a User Group:

1. Navigate to **Administer** → **User groups** to view the user groups on your Satellite.
2. Click **New User Group**.

3. On the **User group** tab, specify the name of the new user group and select group members from the list of users. To include a previously-created user group, select the check box next to the name of the group to be added.
4. On the **Roles** tab, select the roles you want to assign to the user group. Alternatively, select the **Administrator** check box to assign all available permissions.
5. Click **Submit** to create the user group.

4.3. CREATING AND MANAGING ROLES

Red Hat Satellite provides a set of predefined roles with permissions sufficient for standard tasks, as listed in [Table 4.1, “Predefined Roles Available in Red Hat Satellite”](#). It is also possible to configure custom roles, and assign one or more permission filters to them. Permission filters define the actions allowed for a certain resource type. Certain Satellite plug-ins create roles automatically.

Table 4.1. Predefined Roles Available in Red Hat Satellite

Role	Permissions Provided by Role ^[a]
Anonymous	The set of permissions that every user is granted, irrespective of any other roles.
Discovery manager	View, provision, edit, and destroy discovered hosts and manage discovery rules.
Discovery reader	View hosts and discovery rules.
Boot disk access	Download the boot disk.
Red Hat Access Logs	View the log viewer and the logs.
Manager	A most extensive set of permissions, the majority of actions from each resource type is enabled.
Edit partition tables	View, create, edit and destroy partition tables.
View hosts	View hosts.
Edit hosts	View, create, edit, destroy, and build hosts.
Viewer	A passive role that provides the ability to view the configuration of every element of the Satellite structure, logs, and statistics.
Site manager	A restrained version of the Manager role.
Tasks manager	View and edit Satellite tasks.

Role	Permissions Provided by Role ^[a]
Tasks reader	View Satellite tasks.
<p>[a] The exact set of allowed actions associated with predefined roles can be viewed by the privileged user as described in Section 4.3.3, “Viewing Permissions of a Role”.</p>	

4.3.1. Creating a Role

The following steps show how to create a role.

Procedure 4.7. To Create a Role:

1. Navigate to **Administer** → **Roles**.
2. Click **New Role**.
3. Provide a **Name** for the role.
4. Click **Submit** to save your new role.

To serve its purpose, a role must contain permissions. After creating a role, proceed to [Section 4.3.2, “Adding Permissions to a Role”](#).



NOTE

Cloning an existing role is a time-saving method of role creation, especially if you want to create a new role that is a variation of an existing permission set. To clone a role, navigate to **Administer** → **Roles** and select **Clone** from the drop-down list to the right of the role to be copied. Select the name for the new role and alter the permissions as needed.

4.3.2. Adding Permissions to a Role

The following steps show how to add permissions to a role.

Procedure 4.8. To Add Permissions to a Role:

1. Navigate to **Administer** → **Roles**.
2. Select **Add Filter** from the drop-down list to the right of the required role.
3. Select the **Resource type** from the drop-down list. The *(Miscellaneous)* group gathers permissions that are not associated with any resource group.
4. Click the permissions you want to select from the **Permission** list.
5. Select whether the permission is **Unlimited**. This option is selected by default, which means that the permission is applied on all resources of the selected type. When you disable the **Unlimited** check box, the **Search** field activates. In this field you can specify further filtering with use of the Red Hat Satellite 6 search syntax. See [Section 4.4, “Granular Permission Filtering”](#) for details.

6. Click **Next**.
7. Click **Submit** to save changes.

4.3.3. Viewing Permissions of a Role

The following procedure shows how to view permissions assigned to an existing role.

Procedure 4.9. To View Permissions Associated with a Role:

1. Navigate to **Administer** → **Roles**.
2. Click **Filters** to the right of the required role to get to the **Filters** page.

The **Filters** page contains a table of permissions assigned to a role grouped by the resource type. It is also possible to generate a complete table of permissions and actions that you can use on your Satellite system. See [Procedure 4.10, “To Create a Complete Permission Table:”](#) for instructions.

Procedure 4.10. To Create a Complete Permission Table:

1. Ensure that the required packages are installed. Execute the following command on the Satellite Server:

```
# yum install tfm-rubygem-foreman*
```

2. Start the Satellite console with the following command:

```
# foreman-rake console
```

Insert the following code into the console:

```
f = File.open('/tmp/table.html', 'w')

result = Foreman::AccessControl.permissions.sort {|a,b|
  a.security_block <=> b.security_block}.collect do |p|
  actions = p.actions.collect { |a| "<li>#{a}</li>" }
  "<tr><td>#{p.name}</td><td><ul>#{actions.join('')}</ul></td>
<td>#{p.resource_type}</td></tr>"
end.join("\n")

f.write(result)
```

The above syntax creates a table of permissions and saves it to the **/tmp/table.html** file.

3. Press **Ctrl+D** to exit the Satellite console. Insert the following text at the first line of **/tmp/table.html**:

```
<table border="1"><tr><td>Permission name</td><td>Actions</td>
<td>Resource type</td></tr>
```

Append the following text at the end of **/tmp/table.html**:

```
</table>
```

4. Open `/tmp/table.html` in a web browser to view the table.

4.3.4. Removing a Role

The following steps show how to remove an existing role.

Procedure 4.11. To Remove a Role:

1. Navigate to **Administer** → **Roles**.
2. Select **Delete** from the drop-down list to the right of the role to be deleted.
3. In an alert box that appears, click **OK** to delete the role.

4.4. GRANULAR PERMISSION FILTERING

As mentioned in [Section 4.3.2, “Adding Permissions to a Role”](#), Red Hat Satellite provides the ability to limit the configured user permissions to selected instances of a resource type. These granular filters are queries to the Satellite database and are supported by the majority of resource types.

To create a granular filter, specify a query in the **Search** field on the **Edit Filter** page. Deselect the **Unlimited** check box for the field to be active. Queries have the following form:

```
field_name operator value
```

Where:

- *field_name* marks the field to be queried. The range of available field names depends on the resource type. For example, the *Partition Table* resource type offers *family*, *layout*, and *name* as query parameters.
- *operator* specifies the type of comparison between *field_name* and *value*. See [Table 4.2, “Supported Operators for Granular Search”](#) for an overview of applicable operators.
- *value* is the value used for filtering. This can be for example a name of an organization. Two types of wildcard characters are supported: underscore (`_`) provides single character replacement, while percent sign (`%`) replaces zero or more characters.

For most resource types, the **Search** field provides a drop-down list suggesting the available parameters. This list appears after placing the cursor in the search field. For many resource types, it is also possible to combine the queries by using the *and* and *or* operators.

Table 4.2. Supported Operators for Granular Search

Operator	Description
=	<i>Is equal to.</i> An equality comparison that is case-sensitive for text fields.
!=	<i>Is not equal to.</i> An inversion of the = operator.
~	<i>Like.</i> A case-insensitive occurrence search for text fields.
!~	<i>Not like.</i> An inversion of the ~ operator.
^	<i>In.</i> A case-sensitive search for text fields containing a certain string.
!^	<i>Not in.</i> An inversion of the ^ operator.
>, >=	<i>Greater than, greater than or equal to.</i> Supported for numerical fields only.
<, <=	<i>Less than, less than or equal to.</i> Supported for numerical fields only.

For example, the following query applies any permissions specified for the Host resource type only to hosts in the group named host-editors.

```
hostgroup = host-editors
```

You can also limit permissions to a selected environment. To do so, specify the environment name in the **Search** field, for example:

```
Dev
```

As an administrator, you can allow selected users to make changes in a certain part of the environment path. The above filter allows you to work with content while it is in the development stage of the application life cycle, but the content becomes inaccessible once is pushed to production.



NOTE

Satellite does not apply search conditions to create actions. For example, limiting the `create_locations` action with `name = "Default Location"` expression in the search field will not prevent the user from assigning a custom name to the newly created location.

You can limit user permissions to a certain organization or location with use of the permission filter. However, resource types provide a GUI alternative in form of **Locations** and **Organizations** tabs. On these tabs, you can select from the list of available organizations and locations. See [Example 4.1, "Creating an Organization-specific Manager Role"](#).

Example 4.1. Creating an Organization-specific Manager Role

This example shows how to create a manager role restricted to a single organization named `org-1`.

1. Navigate to **Administer** → **Roles**.
2. Clone the existing *Manager* role. Select **Clone** from the drop-down list next to the **Filters** button. You are then prompted to insert a name for the cloned role, for example *org-1 Manager*.
3. Click **Filters** next to *org-1 Manager* to view the filters associated with the role. All filters are marked as unlimited.
4. For each filter, click **Edit**.
5. If the filter contains the **Organizations** tab, navigate to it. Otherwise it is a global setting that cannot be limited.
6. On the **Organizations** tab, select **org-1**. Click **Submit**.
7. The restricted filters are no longer marked as unlimited. Users assigned with the *org-1 Manager* role can now perform management tasks only in the selected organization.

CHAPTER 5. BACKUP AND DISASTER RECOVERY

This chapter describes the minimum backup and restore procedures required to ensure continuity of your Red Hat Satellite deployment and associated data in the event of a disaster. If your deployment uses custom configurations you should take these into account when planning your backup and disaster recovery policy.

5.1. BACKING UP SATELLITE SERVER OR CAPSULE SERVER

This section describes creating a backup of your Satellite Server or Capsule Server and all associated data using the **katello-backup** script. Backing up to a separate location is recommended. Backing up to a separate storage device on a separate system is highly recommended. As no Satellite services will be available during the backup, the backup can be scheduled for a quiet time (for example, using **cron**).

You can also use conventional backup methods such as that described in the [System Backup and Recovery](#) section of the *Red Hat Enterprise Linux 7 System Administrator's Guide*.



NOTE

- When planning a scheduled backup, ensure that no other tasks are scheduled by other administrators for the same time. This is particularly important when administrators are working in different locations and time zones.
- When creating a snapshot or conventional backup, stop all services (Do not do this if using the **katello-backup** script):

```
# katello-service stop
```

Start the services after creating a snapshot or conventional backup:

```
# katello-service start
```

Red Hat Satellite 6.2 uses the **katello-backup** script to make and restore backups. To see the usage statement, enter a command as follows:

```
# katello-backup --help
```

From Satellite 6.2.8, the **katello-backup** creates a time-stamped subdirectory in the backup directory you specify. The **katello-backup** script does not overwrite backups and the correct directory or subdirectory has to be selected when restoring from a backup or an incremental backup. The script will stop and restart services as required.

Known Issues

The **katello-backup** script has the following known issues:

- In Satellite 6.2.7 and earlier:
 - The script does not support using the same directory multiple times for a full backup. Update to Satellite 6.2.8 or newer if possible. Alternatively, use a new directory every time or move the previous backup to a safe location until the new backup has completed successfully.

- If an existing directory is used as the backup target directory, the script will change the group to the **postgres** group. This can have unexpected consequences if other processes are also using the directory. Update to Satellite 6.2.8 or newer if possible. Alternatively, use a new directory or subdirectory exclusively for backups.
- In Satellite 6.2.8 and earlier, using the **--online-backup** option leaves Satellite services down. Update to Satellite 6.2.9 or newer if possible. Additionally, until [Red Hat Bug 1432013](#) is resolved, use only the `/tmp/` or `/var/tmp/` directory to create online backups.
- In Satellite 6.2.9 and earlier, using the **--incremental** option created full backups. This was fixed in [Red Hat Bug 1445989](#) included in the asynchronous erratum [RHBA-2017:1234 - Bug Fix Advisory](#). On Satellite 6.2.9, apply the erratum or upgrade to a newer version if possible.

Checking the Satellite or Capsule Version

To check the version of Satellite 6, use a command as follows:

```
# yum info satellite
```

Red Hat Satellite 6 does not support updating individual packages. To update to the latest minor version, see [Updating Satellite Server, Capsule Server, and Content Hosts](#) in the *Red Hat Satellite 6.2 Installation Guide*. To upgrade to the next major version, see [Upgrading Satellite Server and Capsule Server](#).

Procedure 5.1. To Perform a Full Backup of Satellite Server or Capsule Server:

This procedure performs a full off-line backup. No Satellite services will be available during the backup.

1. Ensure your backup location has enough disk space to contain a copy of the following directories:
 - `/etc/`
 - `/var/lib/pulp/`
 - `/var/lib/mongodb/`
 - `/var/lib/pgsql/`

In Satellite 6.2.7 and earlier, the backup target directory must be a new directory or subdirectory created exclusively for backups. You can use the **du -sh *directory_name*** command to check the space used by a directory.

2. Request other users of Satellite Server or Capsule Server to save any changes and warn them that for the duration of the backup, no Satellite services will be available. Ensure that no other tasks are scheduled for the same time as the backup.
3. Run the backup script:

```
# katello-backup backup_directory
```

The **katello-backup** script stops all services which could impact the backup, performs the backup, then restarts the required services. If the target directory does not exist when trying to create a backup file the script will create it.

**NOTE**

From Satellite 6.2.13 the **katello-backup** script prompts for confirmation that the backup is to proceed. To run the backup without being prompted, add parameter **--assumeyes**.

This process can take a long time to complete, due to the amount of data to copy.

Procedure 5.2. To Perform a Backup without Pulp Content:

This procedure performs an off-line backup but excludes the contents of the Pulp directory. This backup is useful for debugging purposes and is only intended to provide access to configuration files without spending time backing up the Pulp database. You cannot restore from a directory that does not contain Pulp content.

1. Ensure your backup location has enough disk space to contain a copy of the following directories:
 - **/etc/**
 - **/var/lib/mongodb/**
 - **/var/lib/pgsql/**

You can use the **du -sh *directory_name*** command to check the space used by a directory.

2. Request other users of Satellite Server or Capsule Server to save any changes and warn them that for the duration of the backup, no Satellite services will be available. Ensure that no other tasks are scheduled for the same time as the backup.
3. Run the backup script:

```
# katello-backup --skip-pulp-content backup_directory
```

The **katello-backup** script stops all services which could impact the backup, performs the backup, then restarts the required services. If the target directory does not exist when trying to create a backup file the script will create it.

**NOTE**

From Satellite 6.2.13 the **katello-backup** script prompts for confirmation that the backup is to proceed. To run the backup without being prompted, add parameter **--assumeyes**.

Procedure 5.3. To Perform an Incremental Backup:

This procedure performs an off-line backup of any changes since a previous backup. Use a full backup as a reference to make the first incremental backup of a sequence. Keep at least the last known good full backup and a complete sequence of incremental backups to restore from.

1. Ensure your backup location has enough disk space to contain a copy of all changes in the following directories:
 - `/etc/`
 - `/var/lib/pulp/`
 - `/var/lib/mongodb/`
 - `/var/lib/pgsql/`

You can use the `du -sh directory_name` command to check the space used by a directory.

2. Request other users of Satellite Server or Capsule Server to save any changes and warn them that for the duration of the backup, no Satellite services will be available. Ensure that no other tasks are scheduled for the same time as the backup.
3. Run the backup script:

With Pulp content:

```
# katello-backup backup_directory --incremental
backup_directory/previous_time-stamped_subdirectory
```

Without Pulp content:

```
# katello-backup backup_directory --skip-pulp-content --incremental
backup_directory/previous_time-stamped_subdirectory
```

The **katello-backup** script stops all services which could impact the backup, performs the backup, then restarts the required services. If the target directory does not exist when trying to create a backup file the script will create it. It is possible to make incremental backups using a backup older than the previous backup as a starting point, but with a corresponding increase in time to make the backup.




NOTE


From Satellite 6.2.13 the **katello-backup** script prompts for confirmation that the backup is to proceed. To run the backup without being prompted, add parameter **--assumeyes**.

Procedure 5.4. To Perform an Online Backup:

This procedure performs a full backup while Satellite Server or Capsule Server is running. If there are procedures affecting the Pulp database, the Pulp part of the backup procedure will repeat until it is no longer being altered. Since the backup of the Pulp database is the most time consuming part of backing up a Satellite, it is **highly** recommended to not alter the Pulp database during this time. It would prolong the procedure as the Pulp part of the backup will restart.

**NOTE**

Until [Red Hat Bug 1432013](#) is resolved, use only the `/tmp/` or `/var/tmp/` directory to create online backups.

**IMPORTANT**

Satellite 6 uses two database systems, Postgres and Mongo. There are records that exist in both Postgres and Mongo that need to remain synchronized.

The `--online-backup` option keeps all services running which means there is a possibility that data can be modified while the backup is being made. There is a basic check to see if the databases were modified during the backup. If this occurs, the script starts the database portion of the backup again. This check is rudimentary and cannot ensure with 100% certainty that there were no modifications to the databases while the backup script was running. This check can also result in repeated loops if there is continuous modification occurring to the databases.

If you still want to use the `--online-backup` method in production, ensure that no modifications occur during the backup.

1. Ensure you have updated to Satellite 6.2.9 or later. To check the version of Satellite 6, use a command as follows:

```
# yum info satellite
```

2. Ensure your backup location has enough disk space to contain a copy of the following directories:
 - `/etc/`
 - `/var/lib/pulp/`
 - `/var/lib/mongodb/`
 - `/var/lib/pgsql/`

You can use the `du -sh directory_name` command to check the space used by a directory.

3. Request other users of Satellite Server or Capsule Server to save any changes and ask them not to make changes to repositories and Content Views for the duration of the backup. Ensure that no other tasks, such as synchronizing of repositories, are scheduled for the same time as the backup.
4. Run the backup script:

```
# katello-backup --online-backup /tmp/backup_directory
```

**NOTE**

From Satellite 6.2.13 the **katello-backup** script prompts for confirmation that the backup is to proceed. To run the backup without being prompted, add parameter **--assumeyes**.

Example 5.1. A Weekly Full Backup Followed by Daily Incremental Backups

An example script which makes a full backup on a Sunday and incremental backups on all other days of the week:

```
#!/bin/bash -e
DESTINATION=/var/backup
if [[ $(date +%w) == 0 ]]; then
    katello-backup $DESTINATION --assumeyes
else
    LAST=$(ls -td -- $DESTINATION/*/ | head -n 1)
    katello-backup $DESTINATION --incremental "$LAST" --assumeyes
fi
exit 0
```

Add the **--assumeyes** parameter only if your Satellite version is 6.2.13 and higher.

5.2. RESTORING SATELLITE SERVER OR CAPSULE SERVER FROM A BACKUP

This section describes how to restore a Red Hat Satellite Server or Red Hat Capsule Server from the backup data created as a result of following the steps in [Section 5.1, “Backing up Satellite Server or Capsule Server”](#). This process is intended for restoring the backup on the same server that generated the backup, and all data covered by the backup will be deleted on the target system. If the original system is unavailable, provision a system with the same configuration settings (in particular, the host name must be the same).

Prerequisites

Ensure that you address the following conditions:

- Ensure you are restoring to the correct instance. The Red Hat Satellite instance must have the same host name, configuration, and be the same major version as the original system.
- Run the **katello-restore** script as **root**.
- All SELinux contexts must be correct. Enter the following command to restore the correct SELinux contexts:

```
# restorecon -Rnv /
```

Procedure 5.5. To Restore Satellite Server or Capsule Server from a Full Backup:

1. Install Satellite 6 using the procedures in the [Red Hat Satellite 6 Installation Guide](#)^[1].

2. Copy the backup data to the Satellite's local file system. Use `/tmp/` or `/var/tmp/`. Ensure you have enough space to store this data on the base system of Satellite Server or Capsule Server as well as enough space after the restoration to contain all the data in the `/etc/` and `/var/` directories contained within the backup.

You can use the `du -sh directory_name` command to check the space used by a directory and the `df -h directory_name` command to check for free space. Add the `--total` option to sum the results from more than one directory.

3. Run the restoration script:

```
# katello-restore backup_directory
```

Where *backup_directory* is the time-stamped directory or subdirectory containing the backed-up data. The target directory will be read from the configuration files contained within the archive. If the target directory does not exist when trying to recover, it will give an error and ask for the correct directory. The restore process can take a long time to complete, due to the amount of data to copy. Where incremental backups exist, see [Procedure 5.6, "To Restore Satellite Server or Capsule Server from an Incremental Backup:"](#).

When this process completes, all services should be running and Satellite Server or Capsule Server should be available for use.

Procedure 5.6. To Restore Satellite Server or Capsule Server from an Incremental Backup:

1. Install Satellite 6 using the procedures in the [Red Hat Satellite 6 Installation Guide](#)^[2].
2. Restore the last full backup as described in [Procedure 5.5, "To Restore Satellite Server or Capsule Server from a Full Backup:"](#).
3. Copy the backup data to the Satellite's local file system, for example, `/var/tmp/satellite-backup/`. Ensure you have enough space to store this data on the base system of Satellite Server or Capsule Server as well as enough space after the restoration to contain all the data in the `/etc/` and `/var/` directories contained within the backup.
4. Run the restoration script:

```
# katello-restore backup_directory_X
```

Where *backup_directory_X* is a time-stamped directory or subdirectory containing an incremental backup. Restore the incremental backups in the same sequence that they were made. For example: *backup_directory_1*, *backup_directory_2*. The target directory will be read from the configuration files contained within the archive. If the target directory does not exist when trying to recover, it will give an error and ask for the correct directory.

When this process completes, all services should be running and Satellite Server or Capsule Server should be available for use.

5.3. BACKING UP AND RESTORING CAPSULE SERVER USING A SNAPSHOT

There are three methods of backing up Capsule Server:

- Using the **katello-backup** script as described in [Section 5.1, “Backing up Satellite Server or Capsule Server”](#). The **katello-backup** script is convenient if your Capsule Server is a physical machine. You can also use the script if your Capsule Server is a virtual machine but it creates only a backup of the data and not the machine itself.
- Using the conventional backup methods as described in [System Backup and Recovery](#) in the *Red Hat Enterprise Linux 7 System Administrator's Guide*.
- Using the snapshot as described below.

If your Capsule Server is a virtual machine, you can restore it from a snapshot. Creating weekly snapshots to restore from is recommended. In the event of failure, you can reinstall, or configure a new Capsule Server, and then synchronize the database content from the Satellite Server.



NOTE

When creating a snapshot or conventional backup, stop all services (Do not do this if using the **katello-backup** script):

```
# katello-service stop
```

Start the services after creating a snapshot or conventional backup:

```
# katello-service start
```

If you have a snapshot or conventional backup, restore from that and then synchronize from the Satellite Server as described below.

If required, deploy a new Capsule Server, ensuring the host name is the same as before, and then install the Capsule certificates. You might still have them on the Satellite Server, the package name ends in **certs-tar**, alternately create new ones. Follow the procedures in [Installing Capsule Server](#) in the *Red Hat Satellite 6 Installation Guide* until you see in the web UI that the Capsule Server is connected to the Satellite Server. Then synchronize from the Satellite Server as described below.

Procedure 5.7. Synchronizing an External Capsule

1. To synchronize an external Capsule, select the relevant organization and location in the web UI, or choose **Any Organization** and **Any Location**.
2. Navigate to **Infrastructure** → **Capsules** and click the name of the Capsule to synchronize.
3. On the **Overview** tab, select **Synchronize**.

5.4. RENAMING A SATELLITE SERVER OR CAPSULE SERVER

Renaming a Satellite Server or Capsule Server requires use of the **katello-change-**

hostname script. Red Hat Satellite contains references to the host's name and these changes are made using the script. Renaming a Satellite Server affects itself, all Capsule Servers and all hosts registered to it. Renaming a Capsule Server affects itself and all hosts registered to it.



WARNING

The renaming process shuts down all Satellite Server services on the host being renamed. When the renaming is complete, all services are restarted.

5.4.1. Renaming a Satellite Server

The host name of a Satellite Server is used by Satellite Server components, all Capsule Servers, and hosts registered to it for communication. Renaming a Satellite Server requires that these references be updated.

Prerequisites

- (Optional) If the Satellite Server has a custom X.509 certificate installed, a new certificate must be obtained in the host's new name. When all hosts are re-registered to the Satellite Server, the new certificate is installed. For more information on obtaining a custom X.509 certificate, see [Configuring Satellite Server with a Custom Server Certificate](#) in the *Installation Guide*.
- Backup the Satellite Server. The **katello-change-hostname** script makes irreversible changes to the Satellite Server. If the renaming process is not successful, you must restore it from backup. For more information, see [Chapter 5, Backup and Disaster Recovery](#).

Procedure 5.8. Rename a Satellite Server

1. On the Satellite Server, run the **katello-change-hostname** script, providing the host's new name, and Satellite credentials.

```
# katello-change-hostname new_satellite --username admin \  
--password password
```

The message ****** Hostname change complete! ****** confirms that the rename completed successfully.

2. (Optional) If you have obtained a new X.509 certificate for the Satellite Server's new host name, run the Satellite installation script to install the certificate. For more information on installing a custom X.509 certificate, see [Configuring Satellite Server with a Custom Server Certificate](#) in the *Installation Guide*.
3. On all Capsule Servers and hosts registered to the Satellite Server, reinstall the bootstrap RPM and re-register them to the Satellite Server. Substitute the example organization and environment values with those matching your environment.

- a.

```
# yum remove -y katello-ca-consumer*
```
- b.

```
# rpm -Uvh http://new-satellite.example.com/pub/katello-ca-consumer-latest.noarch.rpm
```
- c.

```
# subscription-manager register --org="Default_Organization" \
--environment="Library" \
--force
```

Use of the Red Hat Satellite remote execution feature is recommended for this step. For details, see [Running Jobs on Satellite Hosts](#) in the *Host Configuration Guide*.

4. Reattach subscriptions to all Capsule Servers and hosts registered to the Satellite Server, then refresh the subscription.

- a.

```
# subscription-manager refresh
```
- b.

```
# yum repolist
```

Use of the Red Hat Satellite remote execution feature is recommended for this step. For details, see [Running Jobs on Satellite Hosts](#) in the *Host Configuration Guide*.

5. On all Capsule Servers, re-run the Satellite installation script to update references to the new host name.

```
# satellite-installer --capsule-parent-fqdn new-
satellite.example.com \
--foreman-proxy-foreman-base-url https://new-satellite.example.com
\
--foreman-proxy-trusted-hosts new-satellite.example.com
```

5.4.2. Renaming a Capsule Server

The host name of a Capsule Server is referenced by Satellite Server components, and all hosts registered to it. Renaming a Capsule Server requires that these references be updated.

Prerequisites

- (Optional) New X.509 custom certificate files for the Capsule Server. For more information on obtaining a custom X.509 certificate, see [Configuring Capsule Server with a Custom Server Certificate](#) in the *Installation Guide*.
- Backup the Capsule Server. The **katello-change-hostname** script makes irreversible changes to the Capsule Server. If the renaming process is not successful, you must restore it from backup.

Red Hat Satellite does not provide a native backup method for a Capsule Server. For more information, see [Chapter 5, Backup and Disaster Recovery](#).

Procedure 5.9. Rename a Capsule Server

1. On the Satellite Server, create a new certificates archive file.

- a. If you are using the default Satellite Server certificate:

```
# capsule-certs-generate --capsule-fqdn "new-capsule.example.com" \
--certs-tar "new-capsule.example.com-certs.tar"
```

- b. If you are using a custom X.509 certificate on the Capsule Server, see [Create the Capsule Server's Certificate Archive File](#) in the *Installation Guide*.

2. On the Satellite Server, copy the certificates archive file to the Capsule Server, providing the **root** user's password when prompted. In this example the archive file is copied to the **root** user's home directory, but you may prefer to copy it elsewhere.

```
# scp /root/new-capsule.example.com-certs.tar
root@capsule.example.com:
```

3. On the Capsule Server, run the **katello-change-hostname** script, providing the host's new name, Satellite credentials, and certificates archive filename.

```
# katello-change-hostname new_capsule --username admin \
--password password \
--certs-tar new-capsule.example.com-certs.tar
```

The message ****** Hostname change complete! ****** confirms that the rename completed successfully.

4. (Optional) If you have obtained a new X.509 certificate in the Capsule Server's new host name, run the Satellite installation script to install the certificate. For more information on installing a custom X.509 certificate, see [Configuring Satellite Server with a Custom Server Certificate](#) in the *Installation Guide*.
5. On all hosts registered to the Capsule Server, reinstall the bootstrap RPM and re-register them to the Capsule Server. Substitute the example organization and environment values with those matching your environment.

- a.

```
# yum remove -y katello-ca-consumer*
```

- b.

```
# rpm -Uvh http://new-capsule.example.com/pub/katello-ca-consumer-latest.noarch.rpm
```

- c.

```
# subscription-manager register --org="Default_Organization" \
--environment="Library" \
--force
```

Use of the Red Hat Satellite remote execution feature is recommended for this step. For details, see [Running Jobs on Satellite Hosts](#) in the *Host Configuration Guide*.

6. Reattach subscriptions to all hosts registered to the Capsule Server, then refresh the subscription.

- a. `# subscription-manager refresh`
- b. `# yum repolist`

Use of the Red Hat Satellite remote execution feature is recommended for this step. For details, see [Running Jobs on Satellite Hosts](#) in the *Host Configuration Guide*.

7. Edit the Capsule Server's name.

- a. In the Satellite web UI, navigate to **Infrastructure** → **Capsules**.
- b. Find the Capsule Server in the list, and click **Edit** in that row.
- c. Edit the **Name** and **URL** fields to match the Capsule Server's new host name, then click **Submit**.

8. On your DNS server, add a record for the Capsule Server's new host name, and delete the record for the previous host name.

[1] <https://access.redhat.com/documentation/en/red-hat-satellite/6.2/single/installation-guide/>

[2] <https://access.redhat.com/documentation/en/red-hat-satellite/6.2/single/installation-guide/>

CHAPTER 6. MAINTAINING A RED HAT SATELLITE SERVER

This chapter provides information on how to maintain a Red Hat Satellite Server, including information on relevant log files, how to enable debug logging, how to open a support case and attach the relevant log tar files, and how to use Red Hat Insights to proactively diagnose systems.

6.1. LOGGING AND REPORTING

Red Hat Satellite provides system information in the form of notifications and log files.

Table 6.1. Log Files for Reporting and Troubleshooting

Log File	Description of Log File Content
<code>/var/log/candlepin</code>	Subscription management
<code>/var/log/foreman</code>	Foreman
<code>/var/log/foreman-proxy</code>	Foreman proxy
<code>/var/log/httpd</code>	Apache HTTP server
<code>/var/log/foreman-installer/satellite</code>	Satellite installer
<code>/var/log/foreman-installer/capsule</code>	Capsule Server installer
<code>/var/log/libvirt</code>	Virtualization API
<code>/var/log/mongodb</code>	Satellite database
<code>/var/log/pulp</code>	Celerybeat and Celery startup request messages. After startup is complete, messages are logged to <code>/var/log/messages</code> .
<code>/var/log/puppet</code>	Configuration management
<code>/var/log/rhsm</code>	Subscription management
<code>/var/log/tomcat6</code> and <code>/var/log/tomcat</code>	Apache web server messages for Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7, respectively.
<code>/var/log/messages</code>	Various other log messages related to pulp, rhsm, and goferd.

You can also use the `foreman-tail` command to follow many of the log files related to Satellite. You can run `foreman-tail -l` to list the processes and services that it follows.

On Red Hat Enterprise Linux 7, you can use the *journal* for more extensive logging information. See [Using the Journal](#)^[3] for more information.

6.2. ENABLING DEBUG LOGGING

This section describes how to enable *debug logging* to provide detailed debugging information for Satellite 6.2. Debug logging provides the most detailed log information and can help with troubleshooting issues that can arise with Satellite 6.2 and its components. It is also possible to enable or disable individual loggers for selective logging.

To enable debug logging, modify the `/etc/foreman/settings.yaml` file.

1. Set the Logging Level to "debug"

By default, the logging level is set to **info**, as in the following:

```
:logging:
  :level: info
```

Alter these lines so that they look like this:

```
:logging:
  :level: debug
```



WARNING

Until [BZ#1325939](#) is resolved, setting the sql logging level to **error** is not recommended.

2. Select Individual Logging Types

By default, the end of `/etc/foreman/settings.yaml` looks like this:

```
# Individual logging types can be toggled on/off here
:loggers:
```

Alter the `/etc/foreman/settings.yaml` file so that it looks like this:

```
:loggers:
  :ldap:
    :enabled: true
  :permissions:
    :enabled: true
  :sql:
    :enabled: true
```

3. Restart Katello Services

```
# katello-service restart
```

Complete List of Loggers with their Default Values

```

:app:
  :enabled: true
:ldap:
  :enabled: false
:permissions:
  :enabled: false
:sql:
  :enabled: false

```

6.3. COLLECTING INFORMATION FROM LOG FILES

There are two utilities available to collect information from log files.

Table 6.2. Log Collecting Utilities

Command	Description
foreman-debug	<p>The foreman-debug command collects configuration and log file data for Red Hat Satellite, its back-end services, and system information. This information is collected and written to a tar file. By default, the output tar file is located at /tmp/foreman-debug-xxx.tar.xz.</p> <p>Additionally, the foreman-debug command exports tasks run during the last 60 days. By default, the output tar file is located at /tmp/task-export-xxx.tar.xz. If the file is missing, see the file /tmp/task-export.log to learn why task export was unsuccessful.</p> <p>For more information, run foreman-debug --help.</p> <p>There is no timeout when running this command.</p>
sosreport	<p>The sosreport command is a tool that collects configuration and diagnostic information from a Red Hat Enterprise Linux system, such as the running kernel version, loaded modules, and system and service configuration files. The command also runs external programs (for example: foreman-debug -g) to collect Satellite-specific information, and stores this output in a tar file.</p> <p>By default, the output tar file is located at /var/tmp/sosreport-XXX-20171002230919.tar.xz. For more information, run sosreport --help or see https://access.redhat.com/solutions/3592: <i>What is a sosreport and how can I create one?</i></p> <p>The sosreport command calls the foreman-debug -g and times out after 500 seconds. If your Satellite Server has large log files or many Satellite tasks, support engineers may require the output of sosreport and foreman-debug when you open a support case.</p>



IMPORTANT

Both **foreman-debug** and **sosreport** remove security information such as passwords, tokens, and keys while collecting information. However, the tar files can still contain sensitive information about the Red Hat Satellite Server. Red Hat recommends that you send this information directly to the intended recipient and not to a public target.

6.4. USING LOG FILES IN SUPPORT CASES

You can use the log files and other information described in this chapter to do your own troubleshooting, or you can capture these and many more files, as well as diagnostic and configuration information, to send to Red Hat Support if you need further assistance.

There are two methods to open a support case with Red Hat Support. You can open a support case directly from the Satellite web UI or from the Customer Portal.

- [Section 6.5.5, “Creating Support Cases Using the Red Hat Access Plug-in”](#): How to open a support case from the Satellite web UI
- <https://access.redhat.com/articles/38363>: *How to open and manage a support case on the Customer Portal*

6.5. ACCESSING CUSTOMER PORTAL SERVICES FROM RED HAT SATELLITE

The Red Hat Access pre-installed plug-in lets you access several Red Hat Customer Portal services from within the Satellite web UI.

The Red Hat Access plug-in provides the following services:

- **Search:** Search solutions in the Customer Portal from within the Satellite web UI.
- **Logs:** Send specific parts (snippets) of the log files to assist in problem solving. Send these log snippets to the Red Hat Customer Portal diagnostic tool chain.
- **Support:** Access your open support cases, modify an open support case and open a new support case from within the Satellite web UI.



NOTE

To access Red Hat Customer Portal resources, you must log in with your Red Hat Customer Portal user identification and password.

6.5.1. Searching for Solutions in the Red Hat Access Plug-in

The Red Hat Access plug-in provides search capabilities that look through the solutions database available in the Red Hat Customer Portal.

Procedure 6.1. To Search for Solutions from the Red Hat Satellite Server:

1. In the upper right, click **Red Hat Access** → **Search**.

2. If necessary, log in to the Red Hat Customer Portal. In the main panel on the upper right, click Log In.

**NOTE**

To access Red Hat Customer Portal resources, you must log in with your Red Hat Customer Portal user identification and password.

3. In the **Red Hat Search** field, enter your search query. Search results display in the left-hand **Recommendations** list.
4. In the **Recommendations** list, click a solution. The solution article displays in the main panel.

6.5.2. Using Logs in the Red Hat Access Plug-in

The log file viewer lets you view the log files and isolate log snippets. You can also send the log snippets through the Customer Portal diagnostic tool chain to assist with problem solving.

Procedure 6.2. To Use the Logs Diagnostic Tool from the Red Hat Satellite Server:

1. In the upper right, click **Red Hat Access → Logs**.
2. If necessary, log in to the Red Hat Customer Portal. In the main panel on the upper right, click **Log In**.

**NOTE**

To access Red Hat Customer Portal resources, you must log in with your Red Hat Customer Portal user identification and password.

3. In the left file tree, select a log file and click the file name.
4. Click **Select File**. A pop-up window displays the log file contents.
5. In the log file, highlight any text sections you want diagnosed. The **Red Hat Diagnose** button displays.
6. Click **Red Hat Diagnose**. The system sends the highlighted information to the Red Hat Customer Portal, and provides solutions that closely match the provided log information.
7. If a solution does the following:
 - Matches the problem, click the solution and follow the required steps to troubleshoot the issue.
 - Does not match the problem, click **Open a New Support Case**. The support case is populated with the highlighted text from the log file. See [Section 6.5.5, “Creating Support Cases Using the Red Hat Access Plug-in”](#).

6.5.3. Viewing Existing Support Cases Using the Red Hat Access Plug-in

You can view your existing support case from your Red Hat Satellite Server using the Red Hat Access Plug-in.

Procedure 6.3. To View Existing Support Cases from the Red Hat Satellite Server:

1. In the upper right, click **Red Hat Access** → **Support** → **My Cases**.
2. If necessary, log in to the Red Hat Customer Portal. In the main panel on the upper right, click **Log In**.



NOTE

To access Red Hat Customer Portal resources, you must log in with your Red Hat Customer Portal user identification and password.

3. To search for a specific support case from existing cases, do any of the following:
 1. In the **Search** field, provide a key word or phrase.
 2. From the drop-down list, choose a specific **Case Group**. Your organization has defined **Case Groups** inside the Red Hat Customer Portal.
 3. Choose a Case Status.
4. From the results, choose a specific support case and click the **Case ID**. The support case is ready to view.

6.5.4. Modifying Support Cases Using the Red Hat Access Plug-in

You can update your existing support cases from your Red Hat Satellite Server using the Red Hat Access Plug-in.

Procedure 6.4. To Update Support Cases from the Red Hat Satellite Server Web UI:

1. Complete the instructions from [Section 6.5.3, “Viewing Existing Support Cases Using the Red Hat Access Plug-in”](#)
2. In the support case, scroll down to the marked sections to do the following:
 - **Attachments:** - Attach a local file from the system. Add a file name to make it easier to identify.



NOTE

File names must be less than 80 characters and the maximum file size for attachments uploaded using the web UI is 250 MB. Use FTP for larger files.

- **Case Discussion:** - Add any updated information about the case you wish to discuss with Global Support Services. After adding information, click **Add Comment**.

6.5.5. Creating Support Cases Using the Red Hat Access Plug-in

You can create a new support case from your Red Hat Satellite Server using the Red Hat Access Plug-in.

Procedure 6.5. To Create a New Support Case Using the Red Hat Satellite Server:

1. In the upper right, click **Red Hat Access** → **Support** → **New Case**.
2. If necessary, log in to the Red Hat Customer Portal. In the main panel on the upper right, click Log In.



NOTE

To access Red Hat Customer Portal resources, you must log in with your Red Hat Customer Portal user identification and password.

3. The **Product** and **Product Version** fields are automatically populated. Complete the other relevant fields, as follows:
 - **Summary** — Provide a brief summary of the issue.
 - **Description** — Write a detailed description of the issue.

Based on the summary provided, recommendations for possible solutions display in the main panel.

4. Click **Next**.
5. Choose the appropriate options, as follows:
 - **Severity** — Select the ticket urgency as 4 (low), 3 (normal), 2 (high), or 1 (urgent).
 - **Case Group** — Based on who needs to be notified, create case groups associated with the support case. Select Case Groups in Red Hat Satellite. Create Case Groups within the Customer Portal.
6. Attach the output of **sosreport** and any required files. Add a file description and click **Attach**.



NOTE

- If you have large log files or many Satellite tasks, it is recommended to also attach the output of **foreman-debug**.
- File names must be less than 80 characters and the maximum file size for attachments uploaded using the web UI is 250 MB. Use FTP for larger files.

7. Click **Submit**. The system uploads the case to the Customer Portal, and provides a case number for your reference.

The Red Hat Knowledgebase article <https://access.redhat.com/articles/445443>: *Red Hat Access: Red Hat Support Tool* has additional information, examples, and video tutorials.

6.6. USING RED HAT INSIGHTS WITH SATELLITE SERVER

Red Hat Insights enables you to proactively diagnose systems and downtime related to security exploits, performance degradation and stability failures. You can use the dashboard to quickly identify key risks to stability, security, or performance. You can sort by category, view details of the impact and resolution, and then determine what systems are affected.

Red Hat Insights is installed by default on Satellite Server. Before using Insights with Satellite Server, go to [Red Hat Insights](#) and click **Satellite 6** for the pre-installation checks and to register your Satellite Servers.

6.7. MONITORING SATELLITE SERVER IN THE WEB UI

From the **About** page in the Satellite Server web UI, you can find an overview of the following:

- System Status, including Capsules, Available Providers, Compute Resources, and Plug-ins
- Support information
- System Information
- Backend System Status
- Installed packages

To navigate to the **About** page:

- In the upper right corner of the Satellite Server web UI, click **Administer** → **About**.



NOTE

After Pulp failure, the status of Pulp might show **OK** instead of **Error** for up to 10 minutes due to synchronization delay.

[3] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System_Administrators_Guide/s1-Using_the_Journal.html

CHAPTER 7. MONITORING CAPSULE SERVERS

The following section shows how to use the Satellite web UI to find Capsule information valuable for maintenance and troubleshooting.

7.1. VIEWING GENERAL CAPSULE INFORMATION

Navigate to **Infrastructure** → **Capsules** to view a table of Capsule Servers registered to the Satellite Server. The information contained in the table answers the following questions:

Is the Capsule Server running?

This is indicated by a green icon in the **Status** column. A red icon indicates an inactive Capsule, use the **service foreman-proxy restart** command on the Capsule Server to activate it.

What services are enabled on the Capsule Server?

In the **Features** column you can verify if the Capsule for example provides a DHCP service or acts as a Pulp node. Capsule features can be enabled during installation or configured in addition, see the [Red Hat Satellite Installation Guide](#) for details.

What organizations and locations is the Capsule Server assigned to?

A Capsule server can be assigned to multiple organizations and locations, but only Capsules belonging to the currently selected organization are displayed. To list all Capsules, select **Any Organization** from the context menu in the top left corner.

After changing the Capsule configuration, select **Refresh** from the drop-down menu in the **Actions** column to make sure the Capsule table is up to date.

Click the Capsule name to view further details. At the **Overview** tab, you can find the same information as in the Capsule table. In addition, you can answer to the following questions:

Which hosts are managed by the Capsule Server?

The number of associated hosts is displayed next to the **Hosts managed** label. Click the number to view the details of associated hosts.

How much storage space is available on the Capsule Server?

The amount of storage space occupied by the Pulp content in `/var/lib/pulp`, `/var/lib/pulp/content`, and `/var/lib/mongodb` is displayed. Also the remaining storage space available on the Capsule can be ascertained.

7.2. MONITORING SERVICES

Navigate to **Infrastructure** → **Capsules** and click the name of the selected Capsule. At the **Services** tab, you can find basic information on Capsule services, such as the list of DNS domains, or the number of Pulp workers. The appearance of the page depends on what services are enabled on the Capsule Server. Services providing more detailed status information can have dedicated tabs at the Capsule page (see [Section 7.3, “Monitoring Puppet”](#)).

7.3. MONITORING PUPPET

Navigate to **Infrastructure** → **Capsules** and click the name of the selected Capsule. At the **Puppet** tab you can find the following:

- A summary of Puppet events, an overview of latest Puppet runs, and the synchronization status of associated hosts at the **General** sub-tab.
- A list of Puppet environments at the **Environments** sub-tab.

At the **Puppet CA** tab you can find the following:

- A certificate status overview and the number of autosign entries at the **General** sub-tab.
- A table of CA certificates associated with the Capsule at the **Certificates** sub-tab. Here you can inspect the certificate expiry data, or cancel the certificate by clicking **Revoke**.
- A list of autosign entries at the **Autosign entries** sub-tab. Here you can create an entry by clicking **New** or delete one by clicking **Delete**.

CHAPTER 8. CONFIGURING EXTERNAL AUTHENTICATION

By using external authentication you can derive user and user group permissions from user group membership in an external identity provider. Therefore, you do not have to create these users and maintain their group membership manually on the Satellite Server.

Red Hat Satellite supports four general scenarios for configuring external authentication:

- Using *Lightweight Directory Access Protocol* (LDAP) server as an external identity provider. LDAP is a set of open protocols used to access centrally stored information over a network. For more information, see [Section 8.1, “Using LDAP”](#). Though you can use LDAP to connect to an IdM or AD server, the setup does not support server discovery, cross-forest trusts, or single sign-on with Kerberos on Satellite's web UI.
- Using *Red Hat Enterprise Linux Identity Management* (IdM) server as an external identity provider. IdM deals with the management of individual identities, their credentials and privileges used in a networking environment. For more information see [Section 8.2, “Using Identity Management”](#).
- Using *Active Directory* (AD) integrated with IdM through cross-forest Kerberos trust as an external identity provider. For more information see [Section 8.3.1, “Using Active Directory with Cross-Forest Trust”](#).
- Using direct AD as an external identity provider. For more information see [Section 8.3.2, “Using Active Directory Directly”](#).

The above scenarios are about providing access to the Satellite Server. In addition, hosts provisioned with Satellite can also be integrated with IdM realms. Red Hat Satellite has a realm feature that will automatically manage the life cycle of any system registered to a realm or domain provider. See [Section 8.5, “External Authentication for Provisioned Hosts”](#) for more information.

8.1. USING LDAP

8.1.1. Configure TLS for Secure LDAP (LDAPS)



NOTE

Though direct LDAP integration is covered in this section, Red Hat recommends that you use SSSD and configure it against IdM, AD, or an LDAP server. These preferred configurations are explained elsewhere in this guide.

If you require Red Hat Satellite to use **TLS** to establish a secure LDAP connection (LDAPS), first obtain certificates used by the LDAP server you are connecting to and mark them as trusted on the base operating system of your Satellite Server as described below. If your LDAP server uses a certificate chain with intermediate certificate authorities, all of the root and intermediate certificates in the chain must be trusted, so ensure all certificates are obtained. If you do not require secure LDAP at this time, proceed to [Procedure 8.1, “To Configure LDAP Authentication:”](#).

Obtain the Certificate from the LDAP Server

If you use Active Directory Certificate Services, export the Enterprise PKI CA Certificate using the Base-64 encoded X.509 format. See [How to configure Active Directory](#)

[authentication with TLS on Satellite 6.2](#) for information on creating and exporting a CA certificate from an Active Directory server.

Download the LDAP server certificate to a temporary location on the Red Hat Enterprise Linux system where the Satellite Server is installed and remove it when finished. For example, `/tmp/example.crt`. The filename extensions `.cer` and `.crt` are only conventions and can refer to DER binary or PEM ASCII format certificates.

Trust the Certificate from the LDAP Server

Red Hat Satellite Server requires the CA certificates for LDAP authentication to be individual files in `/etc/pki/tls/certs/` directory.

Use the `install` command to install the imported certificate into the `/etc/pki/tls/certs/` directory with the correct permissions.

```
# install /tmp/example.crt /etc/pki/tls/certs/
```

Enter the following command as `root` to trust the `example.crt` certificate obtained from the LDAP server:

```
# ln -s example.crt /etc/pki/tls/certs/$(openssl x509 -noout -hash -in /etc/pki/tls/certs/example.crt).0
```

Restart the `httpd` service:

- On Red Hat Enterprise Linux 6:

```
# service httpd restart
```

- On Red Hat Enterprise Linux 7:

```
# systemctl restart httpd
```

8.1.2. Configuring Red Hat Satellite to Use LDAP

Follow this procedure to configure LDAP authentication using the web UI. Note that if you need single sign-on functionality with Kerberos on Satellite's web UI, you should use IdM and AD external authentication instead. See [Section 8.2, “Using Identity Management”](#) or [Section 8.3, “Using Active Directory”](#) for more information on those options.

Procedure 8.1. To Configure LDAP Authentication:

1. Set the allow Network Information System (NIS) service boolean to true to prevent SELinux from stopping outgoing LDAP connections:

- For Red Hat Enterprise Linux 6:

```
# setsebool -P allow_ybind on
```

- For Red Hat Enterprise Linux 7:

```
# setsebool -P nis_enabled on
```

2. Navigate to **Administer** → **LDAP Authentication**.

3. Click **New authentication source**.
4. On the **LDAP server** tab, enter the LDAP server's name, host name, port, and server type. The default port is 389, the default server type is POSIX (alternatively you can select FreeIPA or Active Directory depending on the type of authentication server). For **TLS** encrypted connections, select the **LDAPS** check box to enable encryption. The port should change to 636, which is the default for LDAPS.
5. On the **Account** tab, enter the account information and domain name details. See [Section 8.1.3, "LDAP Setting Descriptions and Examples"](#) for descriptions and examples.
6. On the **Attribute mappings** tab, map LDAP attributes to Satellite attributes. You can map Login name, First name, Surname, Email address, and Photo attributes. See [Section 8.1.3, "LDAP Setting Descriptions and Examples"](#) for examples.
7. On the **Locations** tab, select locations from the left table. Selected locations will be assigned to users created from the LDAP authentication source, and available after their first login.
8. On the **Organizations** tab, select organizations from the left table. Selected organizations will be assigned to users created from the LDAP authentication source, and available after their first login.
9. Click **Submit**.

The Satellite Server is now configured to use the LDAP server. If you did not select **Automatically create accounts in Satellite**, see [Section 4.1.1, "Creating a User"](#) to create user accounts manually. If you selected the option, LDAP users can now log in to Satellite using their LDAP accounts and passwords. After they log in for the first time, Satellite administrator will need to assign roles manually. See [Section 4.1.3, "Assigning Roles to a User"](#) to assign user accounts appropriate roles in Satellite.

8.1.3. LDAP Setting Descriptions and Examples

The following table provides a description for each setting in the **Account** tab.

Table 8.1. Account Tab Settings

Setting	Description
Account username	<p>The LDAP user who has read access to the LDAP server. User name is not required if the server allows anonymous reading, otherwise use the full path to the user's object. For example:</p> <pre>uid=\$login,cn=users,cn=accounts,dc=example,dc=com</pre> <p>The \$login variable stores the username entered on the login page as a literal string. The value is accessed when the variable is expanded.</p> <p>The variable cannot be used with external user groups from an LDAP source because Satellite needs to retrieve the group list without the user logging in. Use either an anonymous, or dedicated service user.</p>

Setting	Description
Account password	The LDAP password for the user defined in the Account username field. This field can remain blank if the Account username is using the \$login variable.
Base DN	The top level domain name of the LDAP directory.
Groups base DN	The top level domain name of the LDAP directory tree that contains groups.
LDAP filter	A filter to restrict LDAP queries.
Automatically create accounts in Satellite	If this option is selected, when LDAP users log in to Satellite for the first time, Satellite user accounts are created automatically for them. Users will see a Permissions Denied warning and need to contact their Satellite administrator to have their user account associated with roles.
Usergroup sync	If this option is selected, the user group membership of a user is automatically synchronized when the user logs in, which ensures the membership is always up to date. If this option is cleared, Satellite relies on a Cron job to regularly synchronize group membership (every 30 minutes by default). See Procedure 8.6, "To Configure an External User Group:" for further context.

The following tables show example settings for different types of LDAP connections. All of the examples below use a dedicated service account called *redhat* that has bind, read, and search permissions on the user and group entries. Note that LDAP attribute names are case sensitive.

Table 8.2. Example Settings for Active Directory LDAP Connection

Setting	Example value
Account username	DOMAIN\redhat
Account password	P@ssword
Base DN	DC=example,DC=COM
Groups Base DN	CN=Users,DC=example,DC=com
Login name attribute	userPrincipalName
First name attribute	givenName
Surname attribute	sn

Setting	Example value
Email address attribute	mail

**NOTE**

userPrincipalName allows the use of whitespace in usernames. The login name attribute **sAMAccountName** (which is not listed in the table above) provides backwards compatibility with legacy Microsoft systems. **sAMAccountName** does not allow the use of whitespace in usernames.

Table 8.3. Example Settings for FreeIPA or Red Hat Identity Management LDAP Connection

Setting	Example value
Account username	uid=redhat,cn=users,cn=accounts,dc=example,dc=com
Base DN	dc=example,dc=com
Groups Base DN	cn=groups,cn=accounts,dc=example,dc=com
Login name attribute	uid
First name attribute	givenName
Surname attribute	sn
Email address attribute	mail

Table 8.4. Example Settings for POSIX (OpenLDAP) LDAP Connection

Setting	Example value
Account username	uid=redhat,ou=users,dc=example,dc=com
Base DN	dc=example,dc=com
Groups Base DN	cn=employee,ou=userclass,dc=example,dc=com
Login name attribute	uid

Setting	Example value
First name attribute	givenName
Surname attribute	sn
Email address attribute	mail

8.2. USING IDENTITY MANAGEMENT

Select from one of the following methods:

- [Section 8.2.1, “Using Identity Management Directly”](#)
- [Section 8.2.2, “Using Identity Management with LDAP Authentication”](#)

8.2.1. Using Identity Management Directly

This section shows how to integrate Red Hat Satellite Server with an IdM server and how to enable host-based access control.

Prerequisites

The Satellite Server has to run on Red Hat Enterprise Linux 7.1 or Red Hat Enterprise Linux 6.6 or later.


The examples in this chapter assume separation between IdM and Satellite configuration. However, if you have administrator privileges for both servers, you can configure IdM as described in [Red Hat Enterprise Linux 7 Linux Domain Identity, Authentication, and Policy Guide](#)^[4].

The base operating system of the Satellite Server must be enrolled in the IdM domain by the IdM administrator of your organization.

Procedure 8.2. To Configure IdM Authentication on the Satellite Server:

1. On the IdM server, create a host entry for the Satellite Server and generate a one-time password, for example:

```
# ipa host-add --random hostname
```



NOTE

The generated one-time password must be used on the client to complete IdM-enrollment.

For more information on host configuration properties, see [Red Hat Enterprise Linux 7 Linux Domain Identity, Authentication, and Policy Guide](#)^[5].

2. Create an HTTP service for the Satellite Server, for example:

```
# ipa service-add servicename/hostname
```

For more information on managing services, see [Red Hat Enterprise Linux 7 Linux Domain Identity, Authentication, and Policy Guide](#)^[6].

3. On the Satellite Server, execute the following command as root to configure IdM-enrollment:

```
# ipa-client-install --password OTP
```

Replace *OTP* with the one-time password provided by the IdM administrator.

4. If the Satellite Server is running on Red Hat Enterprise Linux 7, execute the following command:

```
# subscription-manager repos --enable rhel-7-server-optional-rpms
```

The installer is dependent on packages which, on Red Hat Enterprise Linux 7, are in the optional repository **rhel-7-server-optional-rpms**. On Red Hat Enterprise Linux 6 all necessary packages are in the **base** repository.

5. Execute the following command:

```
# satellite-installer --foreman-ipa-authentication=true
```

This command is not limited to a fresh Satellite installation; you can use it to modify an existing Satellite installation.

6. Restart Katello services:

```
# katello-service restart
```

External users can now log in to Satellite using their IdM credentials. They can now choose to either log in to the Satellite Server directly using their username and password or take advantage of the configured Kerberos single sign on and obtain a ticket on their client machine and be logged in automatically. The two-factor authentication with one-time password (2FA OTP) is also supported. If the user in IdM is configured for 2FA, and the Satellite Server is running on Red Hat Enterprise Linux 7, this user can also authenticate to Satellite with a OTP. Optionally proceed to the next procedure to configure host-based access control (HBAC).

HBAC rules define which machine within the domain an IdM user is allowed to access. You can configure HBAC on the IdM server to prevent selected users from accessing the Satellite Server. With this approach, you can prevent Satellite from creating database entries for users that are not allowed to log in. For more information on HBAC, see the [Red Hat Enterprise Linux 7 Linux Domain Identity, Authentication, and Policy Guide](#)^[7]

Procedure 8.3. To Configure HBAC:

1. Create HBAC service and rule on the IdM server and link them together. The following examples use the PAM service name *satellite-prod*. Execute the following commands on the IdM server:

```
$ ipa hbacsvc-add satellite-prod
$ ipa hbacrule-add allow_satellite_prod
$ ipa hbacrule-add-service allow_satellite_prod --
hbacsvcs=satellite-prod
```

2. Add the user who is to have access to the service `satellite-prod`, and the host name of the Satellite Server:

```
$ ipa hbacrule-add-user allow_satellite_prod --user=username
$ ipa hbacrule-add-host allow_satellite_prod --hosts=the-satellite-
fqdn
```

Alternatively, host groups and user groups can be added to the `allow_satellite_prod` rule.

3. To check the status of the rule, execute:

```
$ ipa hbacrule-find satellite-prod
$ ipa hbactest --user=username --host=the-satellite-fqdn --
service=satellite-prod
```

4. Ensure the `allow_all` rule is disabled on the IdM server. For instructions on how to do so without disrupting other services see the [How to configure HBAC rules in IdM](#) article on the Red Hat Customer Portal^[8].
5. Configure the IdM integration with the Satellite Server as described in [Procedure 8.2, “To Configure IdM Authentication on the Satellite Server:”](#). On the Satellite Server, define the PAM service as root:

```
# satellite-installer --foreman-pam-service=satellite-prod
```

8.2.2. Using Identity Management with LDAP Authentication

To attach Identity Management as an external authentication source with no single sign-on support, see [Section 8.1, “Using LDAP”](#) for more information.

8.3. USING ACTIVE DIRECTORY

Select from one of the following methods:

- [Section 8.3.1, “Using Active Directory with Cross-Forest Trust”](#)
- [Section 8.3.2, “Using Active Directory Directly”](#)
- [Section 8.3.3, “Using Active Directory with LDAP Authentication”](#)

8.3.1. Using Active Directory with Cross-Forest Trust

Kerberos can create *cross-forest trust* that defines a relationship between two otherwise separate domain forests. A domain forest is a hierarchical structure of domains; both AD and IdM constitute a forest. With a trust relationship enabled between AD and IdM, users of

AD can access Linux hosts and services using a single set of credentials. For more information on cross-forest trusts, see [Red Hat Enterprise Linux Windows Integration Guide](#)^[9].

From the Satellite point of view, the configuration process is the same as integration with IdM server without cross-forest trust configured. The Satellite Server has to be enrolled in the IPM domain and integrated as described in [Section 8.2, “Using Identity Management”](#). On the IdM server, the following additional steps are required:

1. To enable the HBAC feature, create an external group and add the AD group to it. Add the new external group to a POSIX group. Use this POSIX group in a HBAC rule.
2. Configure `sssd` to transfer additional attributes of AD users. Add these attributes to the `nss` and `domain` sections in `/etc/sss/sss.conf`. For example:

```
[nss]
user_attributes=+mail, +sn, +givenname

[domain/EXAMPLE]
ldap_user_extra_attrs=mail, sn, givenname
```

8.3.2. Using Active Directory Directly

This section shows how to use direct Active Directory (AD) as an external authentication source for the Satellite Server. Direct AD integration means that the Satellite Server is joined directly to the AD domain where the identity is stored. The recommended setup consists of two steps: first enroll Satellite with AD as described in [Procedure 8.4, “To Enroll the Satellite Server with the AD Server:”](#), then finalize the AD integration with use of GSS-proxy as described in [Procedure 8.5, “To Configure Direct AD Integration with GSS-proxy:”](#)

The traditional process of Kerberos authentication in Apache requires the Apache process to have read access to the keytab file. GSS-Proxy allows you to implement stricter privilege separation for the Apache server by removing access to the keytab file while preserving Kerberos authentication functionality. When using AD as an external authentication source for Satellite, it is recommended to implement GSS-proxy, because the keys in the keytab file are the same as the host keys.



NOTE

The AD integration requires the Red Hat Satellite Server to be deployed on Red Hat Enterprise Linux 7.1.

Perform the following procedures on Red Hat Enterprise Linux that acts as a base operating system for your Satellite Server. For the examples in this section `EXAMPLE.ORG` is the Kerberos realm for the AD domain. By completing the procedures, users that belong to the `EXAMPLE.ORG` realm can log in to the Satellite Server.

Prerequisites

Ensure that GSS-proxy and `nfs-utils` are installed:

```
# yum install gssproxy nfs-utils
```

Procedure 8.4. To Enroll the Satellite Server with the AD Server:

1. Install the required packages:

```
# yum install sssd adcli realmd ipa-python
```

2. Enroll the Satellite Server with the AD server. You may need to have administrator permissions to perform the following command:

```
# realm join -v EXAMPLE.ORG
```

After enrolling Satellite with the AD server, you can configure the direct AD integration with GSS-proxy using the **satellite-installer** command. This can be done for already installed Satellite or during the Satellite installation. Note that the Apache user must not have access to the keytab file. Also take note of the effective user ID of the Apache user (that can be found by executing **id apache**). The following procedure uses the example UID 48.

Procedure 8.5. To Configure Direct AD Integration with GSS-proxy:

1. The **satellite-installer** command is by default set for the IdM integration. Change this setting by creating the **/etc/ipa/default.conf** file with the following content:

```
[global]
server = unused
realm = EXAMPLE.ORG
```

2. Create the **/etc/net-keytab.conf** file with the following content:

```
[global]
workgroup = EXAMPLE
realm = EXAMPLE.ORG
kerberos method = system keytab
security = ads
```

3. Create a keytab file for HTTP using the following command:

```
# KRB5_KTNAME=FILE:/etc/gssproxy/http.keytab net ads keytab add HTTP
-U administrator -d3 -s /etc/net-keytab.conf
```

This command fetches the HTTP service keytab file from the AD server and stores it at **/etc/gssproxy/http.keytab**. Make sure this file is owned by the root user and group:

```
# chown root:root /etc/gssproxy/http.keytab
```

4. Insert the following line at the beginning of the **/etc/krb5.conf** file:

```
includedir /var/lib/sss/pubconf/krb5.include.d/
```

5. Create an empty keytab file at **/etc/httpd/conf/http.keytab**:

```
# touch /etc/httpd/conf/http.keytab
```

6. Execute the following command:

```
# satellite-installer --foreman-ipa-authentication=true
```

7. Start and enable the **gssproxy** service:

```
# systemctl restart gssproxy.service
# systemctl enable gssproxy.service
```

8. Configure the Apache server to use GSS-proxy by creating the **/etc/systemd/system/httpd.service** file with the following content:

```
.include /lib/systemd/system/httpd.service
[Service]
Environment=GSS_USE_PROXY=1
```

Apply changes to the service:

```
# systemctl daemon-reload
```

9. Start and enable the **httpd** service:

```
# systemctl restart httpd.service
```

With a running Apache server, users making HTTP requests against the server are authenticated if the client has a valid Kerberos ticket.

Users can now configure Kerberos SSO in their browsers to be able to log in without filling in access credentials in the Satellite GUI. For more information on configuring the Firefox browser see the [Red Hat Enterprise Linux System-Level Authentication Guide](#). Users of the Internet Explorer browser have to add the Satellite Server to the list of Local Intranet or Trusted sites, and turn on the *Enable Integrated Windows Authentication* setting. See the Internet Explorer documentation for details.

NOTE

With direct AD integration, HBAC through IdM is not available. As an alternative, you can use Group Policy Objects (GPO) that enable administrators to centrally manage policies in AD environments. To ensure correct GPO to PAM service mapping, use the following `sssd` configuration:

```
access_provider = ad
ad_gpo_access_control = enforcing
ad_gpo_map_service = +foreman
```

Here, *foreman* is the PAM service name. For more information on GPOs, please refer to the [Red Hat Enterprise Linux Windows Integration Guide](#)^[10].

8.3.3. Using Active Directory with LDAP Authentication

To attach Active Directory as an external authentication source with no single sign-on support, see [Section 8.1, “Using LDAP”](#) for more information. For an example configuration, see [How to configure Active Directory authentication with TLS on Satellite 6](#)

8.4. CONFIGURING EXTERNAL USER GROUPS

Users authenticated through external sources are automatically created on the Satellite Server the first time they log in. This does not apply to external user groups that must be mapped to user groups created manually in the Satellite GUI. Members of the external user group then automatically become members of the Satellite user group and receive the associated permissions.

Prerequisites

The configuration of external user groups depends on the type of external authentication:

- If using an LDAP source, make sure the LDAP authentication is correctly configured. Navigate to **Administer** → **LDAP Authentication** to view and modify the existing sources. For instructions on how to create an LDAP source, see [Section 8.1, “Using LDAP”](#). Take note of the LDAP group names you want to use.



NOTE

If you are using external user groups from an LDAP source, you cannot use the `$login` variable as a substitute for the account user name. You need to use either an anonymous or dedicated service user.

- If your Satellite is enrolled with the IdM or AD server as described in [Chapter 8, Configuring External Authentication](#), take note of the external group names you want to use. To find the group membership of external users, execute the `id` command on Satellite:

```
# id username
```

Here, *username* is the name of the external group member. Note that Satellite allows you to configure external groups only after at least one external user authenticates for the first time. Also, at least one user must exist in the external authentication source.

Procedure 8.6. To Configure an External User Group:

1. Navigate to **Administer** → **User Groups**. Click **New User Group**.
2. On the **User group** tab, specify the name of the new user group. Do not select any users as they will be added automatically when refreshing the external user group.
3. On the **Roles** tab, select the roles you want to assign to the user group. Alternatively, select the **Administrator** check box to assign all available permissions.
4. On the **External groups** tab, click **Add external user group** and select an authentication source from the **Auth source** drop-down menu.

Specify the exact name of the LDAP or external group in the **Name** field.

5. Click **Submit**.



IMPORTANT

You can set the LDAP source to synchronize user group membership automatically on user login. If this option is not set, LDAP user groups are refreshed automatically through a scheduled task (cron job) synchronizing the LDAP Authentication source (every 30 minutes by default). If the user groups in the LDAP Authentication source change in the lapse of time between scheduled tasks, the user can be assigned to incorrect external user groups. This is corrected automatically when the scheduled task runs. You can also refresh the LDAP source manually by executing **foreman-rake ldap:refresh_usergroups** or by refreshing the external user groups through the web UI.

External user groups based on IdM or AD are refreshed only when a group member logs in to Satellite. It is not possible to alter user membership of external user groups in the Satellite GUI, such changes are overwritten on the next group refresh. To assign additional permissions to an external user, add this user to an internal user group that has no external mapping specified. Then assign the required roles to this group.

8.5. EXTERNAL AUTHENTICATION FOR PROVISIONED HOSTS

This section shows how to configure IdM integration to authenticate provisioned hosts. First configure the Satellite or Capsule Server for IdM realm support, then add hosts to the IdM realm group.

8.5.1. Configuring a Red Hat Satellite Server or Capsule Server for IdM Realm Support

To use IdM for provisioned hosts, first configure the Red Hat Satellite Server or Red Hat Satellite Capsule Server.

Prerequisites

- A Satellite Server is registered to the content delivery network, an independent Capsule Server is registered to the Satellite Server.
- A realm or domain provider such as Red Hat Identity Management is configured.

Procedure 8.7. To configure the Satellite Server or Capsule Server for IdM Realm Support:

1. On the Satellite Server or Capsule Server, install the following packages:

```
# yum install ipa-client foreman-proxy ipa-admintools
```

2. Configure the Satellite Server (or Capsule Server) as an IdM client:

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

3. Create a realm-capsule user and the relevant roles in Red Hat Identity Management on the Satellite Server or Capsule Server:

```
# foreman-prepare-realm admin realm-capsule
```

Running `foreman-prepare-realm` will prepare an IdM server for use with the Capsule Server. It creates a dedicated role with the permissions needed for Satellite, creates a user with that role and retrieves the keytab file. You will need your Identity Management server configuration details on this step.

If the command successfully executes, you should be able to see the following command output:

```
Keytab successfully retrieved and stored in: freeipa.keytab
Realm Proxy User:    realm-capsule
Realm Proxy Keytab:  /root/freeipa.keytab
```

4. Move the `/root/freeipa.keytab` to the `/etc/foreman-proxy` directory and set the ownership settings to the user `foreman-proxy`:

```
# mv /root/freeipa.keytab /etc/foreman-proxy
# chown foreman-proxy:foreman-proxy /etc/foreman-
proxy/freeipa.keytab
```

5. Configure the realm based on whether you are using Satellite Server or Capsule Server:

- If you are using the integrated capsule Server in the Satellite Server, use **satellite-installer** to configure the realm:

```
# satellite-installer --foreman-proxy-realm true \
--foreman-proxy-realm-keytab /etc/foreman-proxy/freeipa.keytab \
--foreman-proxy-realm-principal 'realm-capsule@EXAMPLE.COM' \
--foreman-proxy-realm-provider freeipa
```



NOTE

You can also run these options when you first configure the Red Hat Satellite Server.

- If you are using an independent Capsule Server, use **satellite-installer --scenario-capsule** to configure the realm:

```
# satellite-installer --scenario-capsule --realm true \
--realm-keytab /etc/foreman-proxy/freeipa.keytab \
--realm-principal 'realm-capsule@EXAMPLE.COM' \
--realm-provider freeipa
```

6. Make sure that the most updated versions of the `ca-certificates` package is installed and trust the IdM Certificate Authority:

```
# cp /etc/ipa/ca.crt /etc/pki/ca-trust/source/anchors/ipa.crt
# update-ca-trust enable
# update-ca-trust
```

7. (Optional) If you are configuring IdM on an already existing Satellite Server or Capsule Server, the following steps should also be taken to make sure that the configuration changes take effect:
 - a. Restart the foreman-proxy service:


```
# service foreman-proxy restart
```
 - b. Log in to the Satellite Server and click **Infrastructure** → **Capsules**.
 - c. Click on the drop-down menu on the right-hand side of the Capsule Server you have configured for IdM and choose **Refresh Features**.
8. Finally, create a new realm entry in the Satellite Server user interface:
 - a. Click **Infrastructure** → **Realms** and on the right-hand corner of the main page, click **New Realm**.
 - b. Fill in the fields in the following subtabs:
 1. On the **Realm** subtab, provide the realm name, the type of realm to use and the realm proxy.
 2. On the **Locations** subtab, choose the locations where the new realm is intended for use.
 3. On the **Organizations** subtab, choose the organizations where the new realm is intended for use.
 - c. Click **Submit**.

The Satellite Server or Capsule Server is now ready to provision hosts that automatically register to IdM. The next section will detail the steps on how to automatically add hosts to an IdM host group.

8.5.2. Adding Hosts to an IdM Host Group

Red Hat Enterprise Linux Identity Management (IdM) supports the ability to set up automatic membership rules based on a system's attributes. Red Hat Satellite's realm feature provides administrators with the ability to map the Red Hat Satellite host groups to the IdM parameter "userclass" which allow administrators to configure automembership.

When nested host groups are used, they are sent to the IdM server as they are displayed in the Red Hat Satellite User Interface. For example, "Parent/Child/Child".



NOTE

The Satellite Server or Capsule Server sends updates to the IdM server, however automembership rules are only applied at initial registration.

Procedure 8.8. To Add Hosts to an IdM Host Group:

1. On the IdM server, create a host group:

```
# ipa hostgroup-add hostgroup_name
```

```

Description: hostgroup_description
-----
Added hostgroup "hostgroup_name"
-----
Host-group: hostgroup_name
Description: hostgroup_description

```

Where:

1. *hostgroup_name* is the host group's name.
 2. *hostgroup_description* is the host group's description.
2. Create an automembership rule:

```

# ipa automember-add --type=hostgroup automember_rule
-----
Added automember rule "automember_rule"
-----
Automember Rule: automember_rule

```

Where:

1. ***automember-add*** flags the group as an automember group.
 2. ***--type=hostgroup*** identifies that the target group is a host group, not a user group.
 3. *automember_rule* is the name you wish to identify the automember rule by.
3. Define an automembership condition based on the userclass attribute:

```

# ipa automember-add-condition --key=userclass --type=hostgroup --
inclusive-regex=^webserver hostgroup_name
-----
Added condition(s) to "hostgroup_name"
-----
Automember Rule: automember_rule
Inclusive Regex: userclass=^webserver
-----
Number of conditions added 1
-----

```

Where:

1. ***automember-add-condition*** allows you to add regular expression conditions to identify group members.
2. ***--key=userclass*** specifies the key attribute as userclass.
3. ***--type=hostgroup*** identifies that the target group is a host group, not a user group.
4. ***--inclusive-regex=^webserver*** is a regular expression pattern to identify matching values.

5. *hostgroup_name* is the target host group's name.

When a system is added to the Satellite Server's *hostgroup_name* host group, it will now automatically be added to the Identity Management server's "*hostgroup_name*" host group as well. IdM host groups allow for Host-Based Access Controls (HBAC), sudo policies and other IdM functions.

[4] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/linux-manual.html

[5] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/host-attr.html

[6] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/services.l

[7] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/configurir-host-access.html

[8] <https://access.redhat.com/solutions/67895>

[9] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Windows_Integration_Guide/active-directory-trust.html

[10] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Windows_Integration_Guide/sssd-gpo.html

CHAPTER 9. CUSTOMIZING SATELLITE SERVER

Red Hat Satellite Server can be extended by the addition of user interface plug-ins and by the use of hooks triggered by orchestration and Rails events. Some plug-ins are installed by default but additional plug-ins can be installed as RPM packages from the Red Hat repositories and from upstream. Red Hat supports the API but not upstream plug-ins themselves. Some hooks are provided as RPM packages and more hooks can be created as shell scripts. This enables system administrator's familiar with shell scripts to extend the Satellite's capabilities without having to use Ruby and Rails.

9.1. ADDING ADDITIONAL PLUG-INS

To list the plug-ins available from the configured repositories, enter as **root**:

```
# yum search rubygem-foreman
Loaded plugins: product-id, search-disabled-repos, subscription-manager
===== N/S matched: rubygem-foreman
=====
tfm-rubygem-foreman-redhat_access.noarch : Foreman engine to access Red
Hat knowledge base and manage support cases.
tfm-rubygem-foreman-tasks.noarch : Tasks support for Foreman with Dynflow
integration
tfm-rubygem-foreman_abrt.noarch : Display reports from Automatic Bug
Reporting Tool in Foreman
tfm-rubygem-foreman_bootdisk.noarch : Create boot disks to provision hosts
with Foreman
output truncated
```

To view the currently installed plug-ins, enter as **root**:

```
# yum list installed | grep rubygem-foreman
```

To add a new plug-in, install the package and then restart Foreman. For example, to install the SCAP client plug-in, enter as **root**:

```
# yum install rubygem-foreman_scap_client.noarch
```

Restart the Foreman service for the plug-in to be registered:

```
# touch ~foreman/tmp/restart.txt
```

The Foreman website has additional plug-ins [Popular Plugins](#)^[11].



IMPORTANT

Support will be unable to diagnose or support your Satellite if Foreman hooks have been installed and configured. Use Foreman hooks at your own risk.

Red Hat supports the plug-in API but does not provide support for any specific upstream plug-ins themselves. Foreman hooks can modify workflows in Satellite. Because of this, Red Hat support can ask that you remove all hooks in order to get support from Red Hat.

Foreman hooks cannot be migrated by the Satellite migration process. This means that you must remove them before upgrading and then reinstate them after you have confirmed that your Satellite upgrade is working as expected.

Adding Plug-ins from the Foreman Repository

The Foreman repositories are available at yum.theforeman.org/plugins. Separate repositories are available for each Foreman release, containing plug-ins that are compatible with that particular version. Ensure you install plug-ins compatible with the version of Foreman on your system. To determine the Foreman release in use, enter:

```
$ rpm -q foreman
foreman-1.7.2.53-1.el7sat.noarch
```

Configure the Foreman repository as follows:

```
# /etc/yum.repos.d/foreman-plugins.repo
[foreman-plugins]
name=Foreman plugins
baseurl=http://yum.theforeman.org/plugins/1.10/elX/x86_64/
enabled=1
gpgcheck=0
```

Where *X* is **6** or **7** for Red Hat Enterprise Linux 6 or 7 respectively. Change the version number in the URL to match the Foreman release in use. Note the packages are not currently GPG signed.

1. Find the package for the plug-in with the search function. For example, to search for a plug-in with the word "discovery" in the name:

```
# yum search discovery
```

Alternately check the plug-in documentation for the name of the plug-in.

2. Install the package, for example:

```
# yum install tfm-rubygem-foreman_discovery
```

3. Restart the Foreman service for the plug-in to be registered:

```
# touch ~foreman/tmp/restart.txt
```

9.2. USING FOREMAN HOOKS

Foreman's host orchestration can be extended by means of hooks so that additional tasks can be executed. A Foreman hook enables triggering a script (any executable can be used) when an orchestration event occurs, such as when a host is created or when provisioning of a host has completed. In addition, hooks can be made into standard Rails callbacks for any Foreman object, all with scripts.



NOTE

Foreman hooks can modify workflows in Satellite and therefore you might be requested to remove all hooks in order to get support from Red Hat. Foreman hooks also need to be removed before upgrading, and then reinstated after you have confirmed Satellite is working as expected.

Foreman hooks are provided by the `tfm-rubygem-foreman_hooks` package, which is installed by default. If required, to ensure the package is installed and up to date enter as **root**:

```
# yum install tfm-rubygem-foreman_hooks
Loaded plugins: product-id, search-disabled-repos, subscription-manager
Package tfm-rubygem-foreman_hooks-0.3.9-2.el7sat.noarch already installed
and latest version
Nothing to do
```

Foreman hooks are stored in `/usr/share/foreman/config/hooks/`. A subdirectory must be created for each Foreman object, with further subdirectories created for each event name. A Foreman object can be a host or network interface. The path to the hook is as follows:

```
/usr/share/foreman/config/hooks/object/event/hook_script
```

For example, to create a subdirectory for hooks to be activated after the host has completed its operating system installation, enter a command as follows:

```
# mkdir -p /usr/share/foreman/config/hooks/host/managed/before_provision/
```

If you download a script, and the appropriately named directory has been created already, then use the **install** command as follows to ensure the SELinux context is correct:

```
install hook_script
/usr/share/foreman/config/hooks/object/event/hook_script
```

Alternately, if you created the script directly in the event subdirectory then apply the SELinux context by entering as **root**:

```
# restorecon -RvF /usr/share/foreman/config/hooks
```

The SELinux context is **bin_t** on Red Hat Enterprise Linux 6 and **foreman_hook_t** on Red Hat Enterprise Linux 7. Keep in mind that the script is running confined, therefore some actions might be denied by SELinux. Check for actions denied by SELinux by running **aureport -a** or looking in `/var/log/audit/audit.log`.

For further information on debugging SELinux problems and using the **audit2allow** utility:

- On Red Hat Enterprise Linux 6, see [Fixing Problems](#)^[12].

- On Red Hat Enterprise Linux 7, see [Fixing Problems](#)^[13].

Procedure 9.1. Creating a Foreman Hook to Use the `logger` Command

This hook script creates additional log messages each time Foreman provisions a new server.

1. Create the directory structure on the Satellite Server base system:

```
# mkdir -p
/usr/share/foreman/config/hooks/host/managed/before_provision/
```

2. Create the script as follows:

```
# vi
/usr/share/foreman/config/hooks/host/managed/before_provision/10_logger.sh
#!/bin/bash
logger $1 $2
```

The numeric prefix `10` to the file name `_logger.sh` determines the order of execution for scripts in the same subdirectory. Change this prefix to suit your needs.

3. Change the script owner to **foreman**:

```
# chown foreman:foreman 10_logger.sh
```

4. Change the script permissions to allow execution by the user:

```
# chmod u+x 10_logger.sh
```

5. Ensure the SELinux context is correct on all files in the `/usr/share/foreman/config/hooks` directory:

```
# restorecon -RvF /usr/share/foreman/config/hooks/
```

6. To enable the **foreman** user to use the **logger** command, add the following rule to the `/etc/sudoers` file:

```
# vi /etc/sudoers
foreman ALL=(ALL) NOPASSWD:/usr/bin/logger
```

7. Restart the Foreman service for the hook to be registered:

```
# touch ~foreman/tmp/restart.txt
```

Every Foreman or Rail object can have a hook. See the `/usr/share/app/models/` directory or, to get a full list of available models, enter the following commands:

```
# foreman-rake console
> ActiveRecord::Base.descendants.collect(&:name).collect(&:underscore).sort
```

```
=> ["audited/adapters/active_record/audit", "compute_resource",
    "container",
    output truncated
```

This command output also lists some technical tables which are unlikely to be used with Foreman hooks, for example "active_record" or "habtm". These are most commonly used:

- host
- report

9.2.1. Orchestration Events

Foreman supports orchestration tasks for hosts and network interfaces, referred to as objects, when the object is created, updated, and destroyed. These tasks are shown to the user in the web UI. If they fail, they will automatically trigger a rollback of the action. Orchestration hooks can be given a priority, therefore it is possible to order them before or after built-in orchestration steps (before a DNS record is deployed for example).

To add a hook to an event, use the following event names:

- create
- update
- destroy

9.2.2. Rails Events

For hooks on anything apart from hosts and NICs (which support orchestration, as above) then the standard Rails events can be used. Every event has a "before" and "after" hook and these are the most interesting events provided:

- after_create
- before_create
- after_destroy
- before_destroy

The host object has two additional callbacks that you can use:

- **host/managed/after_build** triggers when a host is put into build mode.
- **host/managed/before_provision** triggers when a host completes the OS install.

For the full list of Rails events, see the Constants section at the bottom of the *Ruby on Rails ActiveRecord::Callbacks*^[14] documentation.

9.2.3. Execution of hooks

Hooks are executed in the context of the Foreman server, so usually under the **foreman** user. The first argument is always the event name, enabling scripts to be symbolically linked into multiple event directories. The second argument is the string representation of the object that was hooked, for example the host name for a host:

```
~foreman/config/hooks/host/managed/create/50_register_system.sh create
foo.example.com
```

A JSON representation of the hook object will be passed in on standard input. This JSON is generated by the v2 API views. A utility to read this with **jgrep** is provided in **examples/hook_functions.sh** and sourcing this utility script will be enough for most users. Otherwise, closing standard input is recommended to prevent the pipe buffer from filling which would block the Foreman thread.

```
echo '{"host":{"name":"foo.example.com"}}' \
| ~foreman/config/hooks/host/managed/create/50_register_system.sh \
  create foo.example.com
```

Every hook within the event directory is executed in alphabetical order. For orchestration hooks, an integer prefix in the hook filename will be used as the priority value, thereby influencing when it is executed in relation to DNS, DHCP, VM creation, and other tasks.

9.2.4. Hook Failures and Rollback

If a hook fails, exits with a non-zero return code, the event is logged. For Rails events, execution of other hooks will continue. For orchestration events, a failure will halt the action and rollback will occur. If another orchestration action fails, the hook might be called again to rollback its action. In that case the first argument will change as appropriate, so it must be obeyed by the script (for example, a "create" hook will be called with "destroy" if it has to be rolled back later).

[11] http://projects.theforeman.org/projects/foreman/wiki/List_of_Plugins

[12] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security-Enhanced_Linux/sect-Security-Enhanced_Linux-Troubleshooting-Fixing_Problems.html

[13] https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/SELinux_Users_and_Administrators_Guide/sect-Security-Enhanced_Linux-Troubleshooting-Fixing_Problems.html

[14] <http://api.rubyonrails.org/classes/ActiveRecord/Callbacks.html>