



Red Hat Satellite 6.3

Administering Red Hat Satellite

A guide to administering Red Hat Satellite.

Red Hat Satellite 6.3 Administering Red Hat Satellite

A guide to administering Red Hat Satellite.

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

Legal Notice

Copyright © 2018 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

This 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	5
1.1. LOGGING IN TO RED HAT SATELLITE	5
1.2. CHANGING THE PASSWORD	7
1.3. RESETTING THE ADMINISTRATIVE USER PASSWORD	7
1.4. SETTING A CUSTOM MESSAGE ON THE LOGIN PAGE	8
1.5. CONFIGURING THE NOVNC CONSOLE	8
CHAPTER 2. STARTING AND STOPPING RED HAT SATELLITE	9
CHAPTER 3. MONITORING RESOURCES	10
3.1. USING THE RED HAT SATELLITE CONTENT DASHBOARD	10
3.1.1. Managing Tasks	13
3.2. MONITORING TRENDS	14
CHAPTER 4. SEARCHING AND BOOKMARKING	15
4.1. BUILDING SEARCH QUERIES	15
4.1.1. Query Syntax	15
4.1.2. Operators	15
4.1.3. Values	16
4.2. USING FREE TEXT SEARCH	17
4.3. MANAGING BOOKMARKS	18
4.3.1. Creating Bookmarks	18
4.3.2. Deleting Bookmarks	18
CHAPTER 5. MANAGING USERS AND ROLES	19
5.1. CREATING AND MANAGING USERS	19
5.1.1. Creating a User	19
5.1.2. Editing a User	20
5.1.3. Assigning Roles to a User	20
5.1.4. Adding SSH Keys to a User	21
5.1.5. Deleting SSH keys from a User	21
5.1.6. Configuring Email Notifications	22
5.1.7. Removing a User	23
5.2. CREATING AND MANAGING USER GROUPS	24
5.2.1. Creating a User Group	24
5.2.2. Removing a User Group	24
5.3. CREATING AND MANAGING ROLES	24
5.3.1. Example User Roles	26
5.3.2. Creating a Role	28
5.3.3. Adding Permissions to a Role	29
5.3.4. Viewing Permissions of a Role	29
5.3.5. Removing a Role	30
5.4. GRANULAR PERMISSION FILTERING	31
CHAPTER 6. MANAGING SECURITY COMPLIANCE	34
6.1. SECURITY CONTENT AUTOMATION PROTOCOL	34
6.1.1. SCAP Content	34
6.1.2. XCCDF Profile	34
6.2. CONFIGURING SCAP CONTENT	34
6.2.1. Importing OpenSCAP Puppet Modules	34
6.2.2. Uploading Extra SCAP Content	35
6.3. MANAGING COMPLIANCE POLICIES	35
6.3.1. Compliance Policy	35

6.3.2. Creating a Policy	35
6.3.3. Viewing a Policy	36
6.3.4. Editing a Policy	36
6.3.5. Deleting a Policy	37
6.3.6. Adding a Policy to a Host	37
6.4. TAILORING FILES	37
6.4.1. Creating a Tailoring File	37
6.4.2. Uploading a Tailoring File	37
6.4.3. Assigning a Tailoring File to a Policy	38
6.5. MONITORING COMPLIANCE	38
6.5.1. Compliance Policy Dashboard	39
6.5.2. Compliance Reports Overview	39
6.5.3. Searching Compliance Reports	40
6.5.4. Viewing a Compliance Report	41
6.5.4.1. Evaluation Characteristics	41
6.5.4.2. Compliance and Scoring	42
6.5.4.3. Rule Overview	42
6.5.4.4. Examining Rule Results	42
6.5.5. Compliance Email Notifications	42
6.6. SPECIFICATIONS SUPPORTED BY OPENS CAP	43
CHAPTER 7. BACKING UP AND RESTORING SATELLITE SERVER AND CAPSULE SERVER	44
7.1. BACKING UP SATELLITE SERVER OR CAPSULE SERVER	44
7.1.1. Estimating the size of a Backup	44
7.1.2. Performing a Full Backup of Satellite Server or Capsule Server	45
7.1.3. Performing a Backup without Pulp Content	46
7.1.4. Performing an Incremental Backup	47
7.1.5. Example - A Weekly Full Backup Followed by Daily Incremental Backups	48
7.1.6. Performing an Online Backup	48
7.1.7. Performing a Snapshot Backup	49
7.2. RESTORING SATELLITE SERVER OR CAPSULE SERVER FROM A BACKUP	50
7.3. BACKING UP AND RESTORING CAPSULE SERVER USING A VIRTUAL MACHINE SNAPSHOT	51
7.4. RENAMING A SATELLITE SERVER OR CAPSULE SERVER	52
7.4.1. Renaming a Satellite Server	53
7.4.2. Renaming a Capsule Server	54
CHAPTER 8. MAINTAINING SATELLITE SERVER	57
8.1. LOGGING AND REPORTING	57
8.2. ENABLING DEBUG LOGGING	58
8.3. COLLECTING INFORMATION FROM LOG FILES	59
8.4. USING LOG FILES IN SUPPORT CASES	59
8.5. CLEANING UNUSED TASKS	60
8.6. RECOVERING FROM A FULL DISK	60
8.7. RECLAIMING DISK SPACE FROM MONGODB	61
8.8. ACCESSING CUSTOMER PORTAL SERVICES FROM RED HAT SATELLITE	62
8.8.1. Searching for Solutions in the Red Hat Access Plug-in	62
8.8.2. Using Logs in the Red Hat Access Plug-in	63
8.8.3. Viewing Existing Support Cases Using the Red Hat Access Plug-in	63
8.8.4. Modifying Support Cases Using the Red Hat Access Plug-in	64
8.8.5. Creating Support Cases Using the Red Hat Access Plug-in	64
8.9. USING RED HAT INSIGHTS WITH SATELLITE SERVER	65
8.10. MONITORING SATELLITE SERVER IN THE WEB UI	65
CHAPTER 9. MONITORING CAPSULE SERVER	67

9.1. VIEWING GENERAL CAPSULE INFORMATION	67
9.2. MONITORING SERVICES	67
9.3. MONITORING PUPPET	67
CHAPTER 10. CONFIGURING EXTERNAL AUTHENTICATION	69
10.1. USING LDAP	69
10.1.1. Configure TLS for Secure LDAP (LDAPS)	69
10.1.2. Configuring Red Hat Satellite to Use LDAP	70
10.1.3. LDAP Setting Descriptions and Examples	71
10.1.3.1. Example LDAP Filters	73
10.2. USING IDENTITY MANAGEMENT	74
10.2.1. Using Identity Management Directly	74
10.2.2. Using Identity Management with LDAP Authentication	76
10.3. USING ACTIVE DIRECTORY	76
10.3.1. Using Active Directory with Cross-Forest Trust	77
10.3.2. Using Active Directory Directly	77
10.3.3. Using Active Directory with LDAP Authentication	80
10.4. CONFIGURING EXTERNAL USER GROUPS	80
10.5. EXTERNAL AUTHENTICATION FOR PROVISIONED HOSTS	81
10.5.1. Configuring a Red Hat Satellite Server or Capsule Server for IdM Realm Support	81
10.5.2. Adding Hosts to an IdM Host Group	83
CHAPTER 11. CUSTOMIZING SATELLITE SERVER	86
11.1. ADDING ADDITIONAL PLUG-INS	86
11.2. USING FOREMAN HOOKS	87
11.2.1. Orchestration Events	90
11.2.2. Rails Events	90
11.2.3. Execution of hooks	90
11.2.4. Hook Failures and Rollback	91
APPENDIX A. SETTINGS PARAMETERS	92

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.

Installing the Katello Root CA Certificate

The first time you log in to Satellite, you may see a warning informing you that you are using the default self-signed certificate and you may 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://satellite.example.com/pub`**
2. Select *katello-server-ca.crt*.
3. Import the certificate into your browser.

To Log in to Satellite:

1. Access the Satellite Server using a web browser pointed to the fully qualified domain name:
`https://satellite.example.com/`

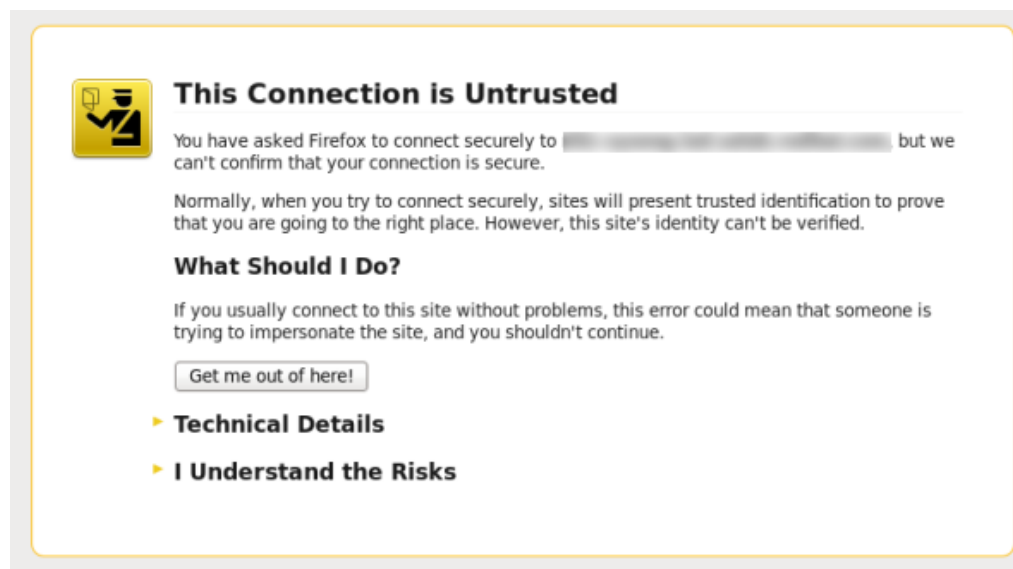
To identify the fully qualified domain name of your Satellite Server, use the **`hostname -f`** command:

```
# hostname -f
```

IMPORTANT

An untrusted connection warning appears in 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.



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”](#).


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](#).

The main navigation tabs are as follows:

Table 1.1. Navigation Tabs

Navigation Tabs	Description
Any Context	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.
Content	Provides content management tools. This includes Content Views, Activation Keys, and Life Cycle Environments.
Containers	Provides container management tools.

Navigation Tabs	Description
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.
Red Hat Access	Provides access to Red Hat knowledgebase, Satellite log files, and support cases.
User Name	Provides user administration where users can edit their personal information.
	Provides event notifications to keep administrators informed of important environment changes.
Administer	Provides advanced configuration for settings such as Users and RBAC, as well as general settings.

1.2. CHANGING THE PASSWORD

These steps show how to change your password.

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 **Current Password** field, enter the current password.
4. In the **Password** field, enter a new password.
5. In the **Verify** field, enter the new password again.
6. Click the **Submit** button to save your new password.

1.3. RESETTING THE ADMINISTRATIVE USER PASSWORD

These steps show how to reset the administrative user password.

To Reset the Administrative User Password:

1. Log in to the machine where Satellite Server is installed.
2. Enter the following command to reset the password:

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

3. Log in to the web user interface and change the password.

1.4. SETTING A CUSTOM MESSAGE ON THE LOGIN PAGE

To Set a Custom Message on the Login Page:

1. Navigate to **Administer > Settings**, and click the **General** tab.
2. Click the edit button next to **Login page footer text**, and enter the desired text to be displayed on the login page. For example, this text may be a warning message required by your company.
3. Click **Save**.
4. Log out of the Satellite's web UI and verify that the custom text is now displayed on the login page below the Satellite version number.

1.5. CONFIGURING THE NOVNC CONSOLE

To access virtual machines using the noVNC console within the Satellite web UI, you must configure network ports and Satellite certificates in your browser. For more information about configuring Satellite to use the NoVNC console, see the [Configuring the NoVNC Console](#) section of the *Red Hat Satellite Provisioning Guide*.

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 creating a backup of Satellite. For more information on creating backups, see [Section 7.1, “Backing up Satellite Server or Capsule Server”](#).

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. MONITORING RESOURCES

The following chapter details how to configure monitoring and reporting for managed systems. This includes host configuration, content views, compliance, subscriptions, registered hosts, promotions and synchronization.

3.1. USING THE RED HAT SATELLITE CONTENT DASHBOARD








The Red Hat Satellite content dashboard contains various widgets which provide an overview of the host configuration, Content Views, compliance reports, subscriptions and hosts currently registered, promotions and synchronization, and a list of the latest notifications.

Navigate to **Monitor > Dashboard** to access the content dashboard. The dashboard can be rearranged by clicking on a widget and dragging it to a different position. The following widgets are available:

Host Configuration Status

An overview of the configuration states and the number of hosts associated with it during the last reporting interval. The following table shows the descriptions of the possible configuration states.

Table 3.1. Host Configuration States

Icon	State	Description
	Hosts that had performed modifications without error	Host that successfully performed modifications during the last reporting interval.
	Hosts in error state	Hosts on which an error was detected during the last reporting interval.
	Good host reports in the last 35 minutes	Hosts without error that did not perform any modifications in the last 35 minutes.
	Hosts that had pending changes	Hosts on which some resources would be applied but Puppet was configured to run in the noop mode.
	Out of sync hosts	Hosts that were not synchronized and the report was not received during the last reporting interval.
	Hosts with no reports	Hosts for which no reports were collected during the last reporting interval.
	Hosts with alerts disabled	Hosts which are not being monitored.

Click the particular configuration status to view hosts associated with it.

Host Configuration Chart

A pie chart shows the proportion of the configuration status and the percentage of all hosts associated with it.

Latest Events

A list of messages produced by hosts including administration information, product and subscription changes, and any errors.

Monitor this section for global notifications sent to all users and to detect any unusual activity or errors.

Run Distribution (last 30 minutes)

A graph shows the distribution of the running Puppet agents during the last puppet interval which is 30 minutes by default. In this case, each column represents a number of reports received from clients during 3 minutes.

New Hosts

A list of the recently created hosts. Click the host for more details.

Task Status

A summary of all current tasks, grouped by their state and result. Click the number to see the list of corresponding tasks.

Latest Warning/Error Tasks

A list of the latest tasks that have been stopped due to a warning or error. Click a task to see more details.

Discovered Hosts

A list of all bare-metal hosts detected on the provisioning network by the Discovery plug-in.

Latest Errata

A list of all errata available for hosts registered to Satellite.

Content Views

A list of all Content Views in Satellite and their publish status.



Sync Overview


An overview of all products or repositories enabled in Satellite and their synchronization status. All products that are in the queue for synchronization, are unsynchronized or have been previously synchronized are listed in this section.

Host Subscription Status

An overview of the subscriptions currently consumed by the hosts registered to Satellite. A subscription is a purchased certificate that unlocks access to software, upgrades, and security fixes for hosts. The following table shows the possible states of subscriptions.

Table 3.2. Host Subscription States

Icon	State	Description
	Invalid	Hosts that have products installed, but are not correctly subscribed. These hosts need attention immediately.
	Partial	Hosts that have a subscription and a valid entitlement, but are not using their full entitlements. These hosts should be monitored to ensure they are configured as expected.

Icon	State	Description
	Valid	Hosts that have a valid entitlement and are using their full entitlements.

Click the subscription type to view hosts associated with subscriptions of the selected type.

Subscription Status

An overview of the current subscription totals that shows the number of active subscriptions, the number of subscriptions that expire in the next 120 days, and the number of subscriptions that have recently expired.

Host Collections

A list of all host collections in Satellite and their status, including the number of content hosts in each host collection.

Virt-who Configuration Status

An overview of the status of reports received from the **virt-who** daemon running on hosts in the environment. The following table shows the possible states.

Table 3.3. Virt-who Configuration States

State	Description
No Reports	No report has been received because either an error occurred during the virt-who configuration deployment, or the configuration has not been deployed yet, or virt-who cannot connect to Foreman during the scheduled interval.
No Change	No report has been received because hypervisor did not detect any changes on the virtual machines, or virt-who failed to upload the reports during the scheduled interval. If you added a virtual machine but the configuration is in the No Change state, check that virt-who is running.
OK	The report has been received without any errors during the scheduled interval.
Total Configurations	A total number of virt-who configurations.

Click the configuration status to see all configurations in this state.

The widget also lists the three latest configurations in the **No Change** state under **Latest Configurations Without Change**.

Latest Compliance Reports

A list of the latest compliance reports. Each compliance report shows a number of rules passed (P), failed (F), or othered (O). Click the host for the detailed compliance report. Click the policy for more details on that policy.

Compliance Reports Breakdown

A pie chart shows the distribution of compliance reports according to their status.

Red Hat Insights Actions

Red Hat Insights is a tool embedded in Satellite that checks the environment and suggests actions you can take. The actions are divided into 4 categories: Availability, Stability, Performance, and Security.

Red Hat Insights Risk Summary

A table shows the distribution of the actions according to the risk levels. Risk level represents how critical the action is and how likely it is to cause an actual issue. The possible risk levels are: Low, Medium, High, and Critical.



NOTE

It is not possible to change the date format displayed in the Satellite web UI.

3.1.1. Managing Tasks

Red Hat Satellite keeps a complete log of all planned or performed tasks, such as repositories synchronised, errata applied, Content Views published, and so on. To review the log, navigate to **Monitor > Tasks**. The page enables you to search for specific tasks, view their status and details, and resume those that resulted in an error, if applicable.

The tasks are managed using the Dynflow engine. Remote tasks have a timeout which can be adjusted as needed.

To Adjust Timeout Settings:

1. Navigate to **Administer > Settings**.
2. Enter `%_timeout` in the search box and click **Search**. The search should return four settings, including a description.
3. In the **Value** column, click the icon next to a number to edit it.
4. Enter the desired value in seconds, and click **Save**.



NOTE

Adjusting the `%_finish_timeout` values might help in case of low bandwidth. Adjusting the `%_accept_timeout` values might help in case of high latency.

When a task is initialized, any back-end service that will be used in the task, such as Candlepin or Pulp, will be checked for correct functioning. If the check fails, you will receive an error similar to the following one:

```
There was an issue with the backend service candlepin: Connection refused
- connect(2).
```

If the back-end service checking feature turns out to be causing any trouble, it can be disabled as follows.

To Disable Checking for Services:

1. Navigate to **Administer > Settings**.
2. Enter `check_services_before_actions` in the search box and click **Search**.
3. In the **Value** column, click the icon to edit the value.
4. From the drop-down menu, select **false**.
5. Click **Save**.

3.2. MONITORING TRENDS


You can use trends to track changes in your infrastructure over time, such as Puppet reports or Facts, and then plan accordingly.

To View a Trend:

1. Navigate to **Monitor > Trends**.
2. On the Trends page, select the trend you want to view from the **Trends** list.

To Create a Trend:

1. Navigate to **Monitor > Trends**.
2. On the Trends page, click the **Add Trend Counter**.
3. From the **Trend type** list, select the category for the new trend.
4. From the **Trendable** list, select the subject for the new trend (if applicable).
5. In the **Name** field, enter a name for the new trend.
6. Click **Submit**.

**NOTE**

If this is the first trend, create a **cron** job to collect trend data:

```
# foreman-rake trends:counter
```

You can set the interval for trend data collection. For example, to collect data once an hour, on the hour:

```
0 * * * * /usr/sbin/foreman-rake trends:counter
```

CHAPTER 4. SEARCHING AND BOOKMARKING

The Satellite web UI features powerful search functionality which is available on most pages of the web UI. It enables you to search all kinds of resources that Satellite Server manages. Searches accept both free text and syntax-based queries, which can be built using extensive input prediction. Search queries can be saved as bookmarks for future reuse.

4.1. BUILDING SEARCH QUERIES

As you start typing a search query, a list of valid options to complete the current part of the query appears. You can either select an option from the list and keep building the query using the prediction, or continue typing. To learn how free text is interpreted by the search engine, see [Section 4.2, “Using Free Text Search”](#).

4.1.1. Query Syntax

parameter operator value

Available fields, resources to search, and the way the query is interpreted all depend on context, that is, the page where you perform the search. For example, the field "hostgroup" on the Hosts page is equivalent to the field "name" on the Host Groups page. The field type also determines available operators and accepted values. For a list of all operators, see [Operators](#). For descriptions of value formats, see [Values](#).

4.1.2. Operators

All operators that can be used between *parameter* and *value* are listed in the following table. Other symbols and special characters that might appear in a prediction-built query, such as colons, do not have special meaning and are treated as free text.

Table 4.1. Comparison Operators Accepted by Search

Operator	Short Name	Description	Example
=	EQUALS	Accepts numerical, temporal, or text values. For text, exact case sensitive matches are returned.	hostgroup = RHEL7
!=	NOT EQUALS		
~	LIKE	Accepts text or temporal values. Returns case insensitive matches. Accepts the following wildcards: <code>_</code> for a single character, <code>%</code> or <code>*</code> for any number of characters including zero. If no wildcard is specified, the string is treated as if surrounded by wildcards: <code>%rhel7%</code>	hostgroup ~ rhel%
!~	NOT LIKE		

Operator	Short Name	Description	Example
>	GREATER THAN	Accepts numerical or temporal values. For temporal values, the operator > is interpreted as "later than", and < as "earlier than". Both operators can be combined with EQUALS: >= <=	registered_at > 10-January-2017 The search will return hosts that have been registered after the given date, that is, between 10th January 2017 and now.
<	LESS THAN		registered_at <= Yesterday The search will return hosts that have been registered yesterday or earlier.
^	IN	Compares an expression against a list of values, as in SQL. Returns matches that contain or not contain the values, respectively.	release_version != 7
!^	NOT IN		
HAS or set?		Returns values that are present or not present, respectively.	has hostgroup or set? hostgroup On the Puppet Classes page, the search will return classes that are assigned to at least one host group. not has hostgroup or null? hostgroup On the Dashboard with an overview of hosts, the search will return all hosts that have no assigned host group.
NOT HAS or null?			

Simple queries that follow the described syntax can be combined into more complex ones using logical operators AND, OR, and NOT. Alternative notations of the operators are also accepted:

Table 4.2. Logical Operators Accepted by Search

Operator	Alternative Notations			Example
and	&	&&	<white space>	class = motd AND environment ~ production
or				errata_status = errata_needed errata_status = security_needed
not	—	!		hostgroup ~ rhel7 not status.failed

4.1.3. Values

Text Values

Text containing whitespaces must be enclosed in quotes. A whitespace is otherwise interpreted as the AND operator.

Examples:

hostgroup = "Web servers"

The search will return hosts with assigned host group named "Web servers".

hostgroup = Web servers

The search will return hosts in the host group Web with any field matching %servers%.

Temporal Values

Many date and time formats are accepted, including the following:

- "10 January 2017"
- "10 Jan 2017"
- 10-January-2017
- 10/January/2017
- "January 10, 2017"
- Today, Yesterday, and the like.



WARNING

Avoid ambiguous date formats, such as 02/10/2017 or 10-02-2017.

4.2. USING FREE TEXT SEARCH

When you enter free text, it will be searched for across multiple fields. For example, if you type "64", the search will return all hosts that have that number in their name, IP address, MAC address, and architecture.



NOTE

Multi-word queries must be enclosed in quotes, otherwise the whitespace is interpreted as the AND operator.

Because of searching across all fields, free text search results are not very accurate and searching can be slow, especially on a large number of hosts. For this reason, we recommend that you avoid free text and use more specific, syntax-based queries whenever possible.

4.3. MANAGING BOOKMARKS

You can save search queries as bookmarks for reuse. You can also delete or modify a bookmark.

Bookmarks appear only on the page on which they were created. On some pages, there are default bookmarks available for the common searches, for example, all **active** or **disabled** hosts.

4.3.1. Creating Bookmarks

This section details how to save a search query as a bookmark. You must save the search query on the relevant page to create a bookmark for that page, for example, saving a host related search query on the Hosts page.

To Create a Bookmark:

1. Navigate to the page where you want to create a bookmark.
2. In the **Search** field, enter the search query you want to save.
3. Select the arrow to the right of the **Search** button and then select **Bookmark this search**.
4. In the **Name** field, enter a name for the new bookmark.
5. In the **Search query** field, ensure your search query is correct.
6. Ensure the **Public** check box is set correctly:
 - Select the **Public** check box to set the bookmark as public and visible to all users.
 - Clear the **Public** check box to set the bookmark as private and only visible to the user who created it.
7. Click **Submit**.

To confirm the creation, either select the arrow to the right of the **Search** button to display the list of bookmarks, or navigate to **Administer > Bookmarks** and then check the **Bookmarks** list for the name of the bookmark.

4.3.2. Deleting Bookmarks

You can delete bookmarks on the Bookmarks page.

To Delete a Bookmark:

1. Navigate to **Administer > Bookmarks**.
2. On the Bookmarks page, click **Delete** for the Bookmark you want to delete.
3. When the confirmation window opens, click **OK** to confirm the deletion.

To confirm the deletion, check the **Bookmarks** list for the name of the bookmark.

CHAPTER 5. MANAGING 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 5.1, “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 5.2, “Creating and Managing 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 5.3, “Creating and Managing Roles”](#).

5.1. 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.

5.1.1. Creating a User

The following steps show how to create a user:

To Create a User:

1. Navigate to **Administer > Users**.
2. Click **Create User**.
3. In the **Username** field, enter the user name for this user to log in to the web UI.
4. In the **First name** and **Surname** fields, enter the real first name and surname of the user.
5. In the **Email address** field, enter the user’s email address.
6. In the **Description** field, insert a description of the new user.
7. Optionally, select a specific language for the user from the **Language** drop-down menu. The default is to attempt to use the language settings of the user’s browser.
8. Optionally, select a specific time zone for the user from the **Timezone** drop-down menu. The default is to use the time zone settings of the user’s browser.
9. Set a password for the user:
 - From the **Authorized by** drop-down menu, select the source by which the user is authenticated. You can select **INTERNAL** to enable the user to be managed inside Satellite Server, or configure an external authentication, such as **LDAP** or **IdM** as described in [Chapter 10, Configuring External Authentication](#).
 - Enter an initial password for the user in the **Password** field and verify it in the **Verify** field.

10. Click **Submit** to create the user.
11. Select the user name to continue configuring the user.
12. On the **Email Preferences** tab, select the **Mail enabled** check box to enable email notifications for the user. Depending on the roles assigned, notifications options can be configured here.
13. On the **Locations** tab, select locations to be made accessible for this user. If you assign multiple locations to the user, you can select the default location for user login from the **Default on login** drop-down menu. Otherwise, when the user logs in, the location selection is set to **Any Location**.
14. On the **Organizations** tab, select organizations to be made accessible to this user. If you assign multiple organizations to the user, you can select the default organization for user login from the **Default on login** drop-down menu. Otherwise, when the user logs in, the organization selection is set to **Any Organization**.
15. On the **Roles** tab, select the required roles for this user.
16. On the **SSH Keys** tab you can add SSH public keys but not until the user has been saved.
17. Click **Submit** to save the changes.

5.1.2. Editing a User

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

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 appears 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.

5.1.3. 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 5.3.2, "Creating a Role"](#). You can apply a similar procedure to user groups.

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 5.1, “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 is removed.

5.1.4. Adding SSH Keys to a User

The following steps show how to add public SSH keys to an existing user. This allows deployment of SSH keys during provisioning.

To deploy SSH keys during provisioning, see [Deploying SSH Keys during Provisioning](#) in the *Red Hat Satellite Provisioning Guide*.

For information on SSH keys and SSH key creation, see [Generating Key Pairs](#) in the *Red Hat Enterprise Linux 7 System Administrator's Guide*.



NOTE

Make sure that you are logged in to the web UI as an Admin user of Red Hat Satellite or a user with the `create_ssh_key` permission enabled.

To Add SSH Keys to a User:

1. Prepare the content of the public SSH key in a clipboard.
2. Navigate to **Administer > Users**.
3. From the **Username** column, click on the username.
4. Select the **SSH Keys** tab.
5. Click **Create SSH Key**.
6. In the **Key** field, paste the content of the public SSH key.
7. In the **Name** field, enter a name for the SSH key.
8. Click **Submit**. A confirmation notification is displayed if your key submission was successful.

5.1.5. Deleting SSH keys from a User

The following steps show how to delete public SSH keys from an existing user.

To Delete SSH Keys from a User:

1. Log in to the Satellite web UI as Admin or as a user with the `destroy_ssh_key` permission enabled.
2. Navigate to **Administer > Users**.
3. Click on the username from the **Username** column.
4. Select the **SSH Keys** tab.
5. Click **Delete** on the row of the SSH key to be deleted.
6. Click **OK** in the confirmation prompt. A confirmation appears to indicate the deletion was successful.

5.1.6. 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*.

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 is associated with the user account, and the notifications selected in the following steps are 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 this triggers 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

5.1.7. Removing a User

The following procedure describes how to remove an existing user.

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.

5.2. CREATING AND MANAGING 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 10.4, “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.

5.2.1. Creating a User Group

The following procedure shows how to create a user group.

To Create a User Group:

1. Navigate to **Administer > User groups**.
2. Click **Create User Group**.
3. On the **User group** tab, specify the name of the new user group and select group members:
 - Select the previously created user groups from the **User Groups** list.
 - Select users from the **Users** list.
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**.

5.2.2. Removing a User Group

The following procedure shows how to remove an existing user group:

To Remove a User Group:

1. Navigate to **Administer > User groups**.
2. Click **Delete** to the right of the user group you want to delete.
3. In the alert box that appears, click **OK** to delete a user group.

5.3. CREATING AND MANAGING ROLES

Red Hat Satellite provides a set of predefined roles with permissions sufficient for standard tasks, as listed in [Table 5.1, “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 5.1. Predefined Roles Available in Red Hat Satellite

Role	Permissions Provided by Role [a].
Access Insights Admin	Add and edit Insights rules.
Access Insights Viewer	View Insight reports.
Boot disk access	Download the boot disk.
Compliance manager	View, create, edit, and destroy SCAP content files, compliance policies, and tailoring files. View compliance reports.
Compliance viewer	View compliance reports.
Create ARF report	Create compliance reports.
Default role	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.
Edit hosts	View, create, edit, destroy, and build hosts.
Edit partition tables	View, create, edit and destroy partition tables.
Manager	A role similar to administrator, but does not have permissions to edit global settings. In the Satellite web UI, global settings can be found under Administer > Settings .
Organization admin	An administrator role defined per organization. The role has no visibility into resources in other organizations.
Red Hat Access Logs	View the log viewer and the logs.
Remote Execution Manager	A role with full remote execution permissions, including modifying job templates.
Remote Execution User	Run remote execution jobs.
Site manager	A restrained version of the Manager role.
Tasks manager	View and edit Satellite tasks.
Tasks reader	A role that can only view Satellite tasks.

Role	Permissions Provided by Role [a].
Viewer	A passive role that provides the ability to view the configuration of every element of the Satellite structure, logs, reports, and statistics.
View hosts	A role that can only view hosts.
Virt-who Manager	A role with full virt-who permissions.
Virt-who Reporter	Upload reports generated by virt-who to Satellite. It can be used if you configure virt-who manually and require a user role that has limited virt-who permissions.
Virt-who Viewer	View virt-who configurations. Users with this role can deploy virt-who instances using existing virt-who configurations.
[a] The exact set of allowed actions associated with predefined roles can be viewed by the privileged user as described in Viewing Permissions of a Role	

5.3.1. Example User Roles

Satellite Administrator

A top level administrator role with access control for all Satellite items, including managed systems and applications.

IT Operations Manager

A read-only role with permissions for viewing Satellite items.

License Management Owner

A task specific role with permissions for managing manifests and subscriptions, including permissions for viewing organizations and reports.

Quality Assurance

An environment and location specific role for testing in a dedicated testing environment but with limited access to items outside that environment.

Table 5.2. Example User Role Configurations

Role	Resource Type	Permissions	Filters
Satellite Administrator	Ensure the Administrator check box is selected for the user. For more information, see To Assign a Role to a User .	predefined permission set	
IT Operations Manager	Viewer	predefined permissions set	

Role	Resource Type	Permissions	Filters
License Management Owner	Miscellaneous	access_dashboard my_organizations view_statistics	
	Products and Repositories	view_products	
	Subscription	view_subscriptions attach_subscriptions unattach_subscriptions import_manifest delete_manifest	
	Organization	view_organizations	
	Report	view_reports	
	Host	view_hosts	
Quality Assurance	Organization	view_organizations	
	Environment	view_environments create_environments edit_environments destroy_environments import_environments	
	Miscellaneous	view_tasks view_statistics access_dashboard	
	Host class	edit_classes	
	Host Group	view_hostgroups edit_hostgroups	

Role	Resource Type	Permissions	Filters
	Host	view_hosts create_hosts edit_hosts destroy_hosts build_hosts power_hosts console_hosts ipmi_boot_hosts puppetrun_hosts	
	Location	view_locations	
	Puppet class	view_puppetclasses	
	Capsule	view_smart_proxies view_smart_proxies_autosign view_smart_proxies_puppetca	
	Miscellaneous	my_organizations	
	Products and Repositories	view_products	
	Host class	edit_classes	
	Lifecycle Environment	view_lifecycle_environments edit_lifecycle_environments promote_or_remove_content_views_to_environments	name ~ QA
	Content Views	view_content_views create_content_views edit_content_views publish_content_views promote_or_remove_content_views	name ~ ccv*

5.3.2. Creating a Role

The following steps show how to create a role.

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 5.3.3](#), “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.

5.3.3. Adding Permissions to a Role

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

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. Depending on the **Resource type** selected, you can select or deselect the **Unlimited** and **Override** check box. The **Unlimited** checkbox 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 5.4](#), “Granular Permission Filtering” for details. When you enable the **Override** check box, you can add additional locations and organizations to allow the role to access the resource type in the additional locations and organizations; you can also remove an already associated location and organization from the resource type to restrict access.
6. Click **Next**.
7. Click **Submit** to save changes.

5.3.4. Viewing Permissions of a Role

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

To View Permissions Associated with a Role:

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 [To Create a Complete Permission Table](#): for instructions.

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.

5.3.5. Removing a Role

The following steps show how to remove an existing role.

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.

5.4. GRANULAR PERMISSION FILTERING

As mentioned in [Section 5.3.3, “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 5.3, “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 5.3. 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> An equality comparison that is case-sensitive search for text fields. This generates a different SQL query to the <i>Is equal to</i> comparison, and is more efficient for multiple value comparison.
!^	<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
```


The following query returns records where the name matches XXXX, Yyyy, or zzzz example strings:

```
name ^ (XXXX, Yyyy, zzzz)
```

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 does 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 the use of the granular permission filter in the **Search** field. However, some resource types provide a GUI alternative, an **Override** check box that provides the **Locations** and **Organizations** tabs. On these tabs, you can select from the list of available organizations and locations. See [Example 5.1, “Creating an Organization-specific Manager Role”](#).

Example 5.1. Creating an Organization-specific Manager Role

This example shows how to create an administrative role restricted to a single organization named *org-1*.

1. Navigate to **Administer > Roles**.
2. Clone the existing **Organization admin** 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 admin*.
3. Click the desired locations and organizations to associate them with the role.
4. Click **Submit** to create the role.
5. Click *org-1 admin*, and click **Filters** to view all associated filters. The default filters work for most use cases. However, you can optionally click **Edit** to change the properties for each filter. For some filters, you can enable the **Override** option if you want the role to be able to



access resources in additional locations and organizations. For example, by selecting the **Domain** resource type, the **Override** option, and then additional locations and organizations using the **Locations** and **Organizations** tabs, you allow this role to access domains in the additional locations and organizations that is not associated with this role. You can also click **New filter** to associate new filters with this role.

CHAPTER 6. MANAGING SECURITY COMPLIANCE

Security compliance management is the ongoing process of defining security policies, auditing for compliance with those policies and resolving instances of non-compliance. Any non-compliance is managed according to the organization's configuration management policies. Security policies range in scope from host-specific to industry-wide, therefore, flexibility in their definition is required.

6.1. SECURITY CONTENT AUTOMATION PROTOCOL

Satellite 6 uses the Security Content Automation Protocol (SCAP) to define security configuration policies. For example, a security policy might specify that for hosts running Red Hat Enterprise Linux, login via SSH is not permitted for the **root** account. With Satellite 6 you can schedule compliance auditing and reporting on all hosts under management. For more information about SCAP, see the [Red Hat Enterprise Linux 7 Security Guide](#).

6.1.1. SCAP Content

SCAP content is a datastream format containing the configuration and security baseline against which hosts are checked. Checklists are described in the extensible checklist configuration description format (XCCDF) and vulnerabilities in the open vulnerability and assessment language (OVAL). Checklist items, also known as rules express the desired configuration of a system item. For example, you may specify that no one can log in to a host over SSH using the **root** user account. Rules can be grouped into one or more profiles, allowing multiple profiles to share a rule. SCAP content consists of both rules and profiles.

You can either create SCAP content or obtain it from a vendor. Supported profiles are provided for Red Hat Enterprise Linux in the `scap-security-guide` package. The creation of SCAP content is outside the scope of this guide, but see the [Red Hat Enterprise Linux 7 Security Guide](#) or [Red Hat Enterprise Linux 6 Security Guide](#) for information on how to download, deploy, modify, and create your own content. The SCAP content provided with Red Hat Enterprise Linux is compliant with SCAP specification 1.2.

The default SCAP content provided with the OpenSCAP components of Satellite 6 depends on the version of Red Hat Enterprise Linux:

- On Red Hat Enterprise Linux 6, content for Red Hat Enterprise Linux 6 is installed.
- On Red Hat Enterprise Linux 7, content for both Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 is installed.

6.1.2. XCCDF Profile

An XCCDF profile is a checklist against which a host or host group is evaluated. Profiles are created to verify compliance with an industry standard or custom standard.

The profiles provided with Satellite 6 are obtained from the [OpenSCAP project](#).

To list available XCCDF profiles, open the Satellite web UI and navigate to **Hosts > SCAP contents**.

6.2. CONFIGURING SCAP CONTENT

6.2.1. Importing OpenSCAP Puppet Modules

To import the OpenSCAP content into a Puppet environment, you must associate each host that you want to audit with the Puppet environment:

1. Navigate to **Configure > Environments**.
2. Click **Import environments from *satellite.example.com***.
3. Select the Puppet environment check box associated with the host you want to audit.
If no Puppet environment exists, select the **production** environment check box.
4. Click **Update**.

6.2.2. Uploading Extra SCAP Content

You can upload extra SCAP content into the Satellite Server, either content created by yourself or obtained elsewhere. SCAP content must be imported into the Satellite Server before being applied in a policy. For example, the **scap-security-guide** RPM package available in the Red Hat Enterprise Linux 7.2 repositories includes a profile for the Payment Card Industry Data Security Standard (PCI-DSS) version 3. You can upload this content into a Satellite Server even if it is not running Red Hat Enterprise Linux 7.2 as the content is not specific to an operating system version.

To Load the Default OpenSCAP Content:

- Load the OpenSCAP content on the Satellite Server using the following command:

```
# foreman-rake foreman_openscap:bulk_upload:default
```

To Upload Extra SCAP Content:

1. Log in to the Satellite web UI.
2. Navigate to **Hosts > SCAP contents** and click **New SCAP Content**.
3. Enter a title in the **Title** text box. For example: **RHEL 7.2 SCAP Content**.
4. Click **Choose file**, navigate to the location containing the SCAP content file and select **Open**.
5. Click **Submit**.

If the SCAP content file is loaded successfully, a message similar to **Successfully created RHEL 7.2 SCAP Content** is shown and the list of **SCAP Contents** includes the new title.

6.3. MANAGING COMPLIANCE POLICIES

6.3.1. Compliance Policy

A scheduled audit, also known as a *compliance policy*, is a scheduled task that checks the specified hosts for compliance against an XCCDF profile. The schedule for scans is specified by the Satellite Server and the scans are performed on the host. When a scan completes, an *Asset Reporting File* (ARF) is generated in XML format and uploaded to the Satellite Server. You can see the results of the scan in the compliance policy dashboard. No changes are made to the scanned host by the compliance policy. The SCAP content includes several profiles with associated rules but policies are not included by default.

6.3.2. Creating a Policy

Follow these steps to create a compliance policy, which specifies the SCAP content and profile to be applied to a location and either a host or host group at a specified time.

Prerequisite

- [Section 6.2.1, “Importing OpenSCAP Puppet Modules”](#).

To Create a Policy:

1. In the Satellite web UI, navigate to **Hosts > Policies**, click **New Policy** and follow the wizard's steps.
2. Enter a name for this policy, a description (optional), then click **Next**.
3. Select the SCAP Content and XCCDF Profile to be applied, then click **Next**.
4. Specify the scheduled time when the policy is to be applied, then click **Next**.
Select **Weekly**, **Monthly**, or **Custom** from the **Period** drop-down list.
 - If you select **Weekly**, also select the desired day of the week from the **Weekday** drop-down list.
 - If you select **Monthly**, also specify the desired day of the month in the **Day of month** field.
 - If you select **Custom**, enter a valid Cron expression in the **Cron line** field.
The **Custom** option allows for greater flexibility in the policy's schedule than either the **Weekly** or **Monthly** options.
5. Select the locations to which the policy is to be applied, then click **Next**.
6. Select the organizations to which the policy is to be applied, then click **Next**.
7. Select the host groups to which the policy is to be applied, then click **Submit**.

When the Puppet agent runs on the hosts which belong to the selected host group, or hosts to which the policy has been applied, the OpenSCAP client will be installed and a Cron job added with the policy's specified schedule. The **SCAP Content** tab provides the name of the SCAP content file which will be distributed to the directory `/var/lib/openscap/content/` on all target hosts.

6.3.3. Viewing a Policy

Follow these steps to preview the rules which will be applied by specific OpenSCAP content and profile combination. This is useful when planning policies.

To View a Policy:

1. In the Satellite web UI, navigate to **Hosts > Policies**.
2. Click **Show Guide**.

6.3.4. Editing a Policy

Follow these steps to edit a policy. An edited policy is applied to the host when its Puppet agent next checks with the Satellite Server for updates. By default this occurs every 30 minutes.

To Edit a Policy:

To Edit a Policy:

1. In the Satellite web UI, navigate to **Hosts > Policies**.
2. From the drop-down list to the right of the policy's name, select **Edit**.
3. Edit the necessary attributes.
4. Click **Submit**.

An edited policy is applied to the host when its Puppet agent next checks with the Satellite Server for updates. By default this occurs every 30 minutes.

6.3.5. Deleting a Policy

Follow these steps to delete an existing policy.

1. In the Satellite web UI, navigate to **Hosts > Policies**.
2. From the drop-down list to the right of the policy's name, select **Delete**.
3. Click **OK** in the confirmation message.

6.3.6. Adding a Policy to a Host

Follow these steps to add a policy to one or more hosts.

1. In the Satellite web UI, navigate to **Hosts > All hosts**.
2. Select the host or hosts to which you want to add the policy.
3. Click **Select Action**.
4. Select **Assign Compliance Policy** from the list.
5. In the new panel that opens, select the appropriate policy from the list of available policies and click **Submit**.

6.4. TAILORING FILES

Tailoring Files allow existing OpenSCAP policies to be customised without forking or rewriting the policy. You can assign a Tailoring File to a policy when creating or updating a policy.

6.4.1. Creating a Tailoring File

You can create a Tailoring File using the [SCAP Workbench](#). For more information on using the SCAP Workbench tool, see [Customizing SCAP Security Guide for your use-case](#).

6.4.2. Uploading a Tailoring File

The following steps show how to upload a Tailoring File:

To Upload a Tailoring File:

1. Log in to the Satellite web UI.

2. Navigate to **Hosts > Compliance - Tailoring Files** and click **New Tailoring File**.
3. Enter a name in the **Name** text box.
4. Click **Choose File**, navigate to the location containing the SCAP DataStream Tailoring File and select **Open**.
5. Click **Submit** to upload the chosen Tailoring File.

6.4.3. Assigning a Tailoring File to a Policy

The following steps show how to assign a Tailoring File to a Policy:

To Assign a Tailoring File to a Policy:

1. Log in to the Satellite web UI.
2. Navigate to **Hosts > Compliance - Policies**.
3. Click **New Policy**, or **New Compliance Policy** if there are existing Compliance Policies.
4. Enter a name in the **Name** text box, and click **Next**.
5. Select a **Scap content** from the dropdown menu.
6. Select a **XCCDF Profile** from the dropdown menu.
7. Select a **Tailoring File** from the dropdown menu.
8. Select a **XCCDF Profile in Tailoring File** from the dropdown menu.
It is important to select the XCCDF Profile because Tailoring Files are able to contain multiple XCCDF Profiles.
9. Click **Next**.
10. Select a **Period** from the dropdown menu.
11. Select a **Weekday** from the dropdown menu, and click **Next**.
12. Select a **Location** to move it to the **Selected Items** window, and click **Next**.
13. Select an **Organization** to move it to the **Selected Items** window, and click **Next**.
14. Select a **Hostgroup** to move it to the **Selected Items** window, and click **Submit**.

6.5. MONITORING COMPLIANCE

Monitoring compliance is an ongoing task of ensuring that audits are conducted and that non-compliance is identified. Red Hat Satellite 6 enables centralized compliance monitoring and management. Hosts under Satellite management are checked for compliance according to your custom schedule and details are collated by the Satellite Server. A compliance dashboard provides an overview of hosts' compliance and the ability to view details for each host within the scope of that policy. Compliance reports provide a detailed analysis of each host's compliance with the applicable policy. With this information you can evaluate the risks presented by each host and better manage the resources required to bring hosts into compliance.

Common objectives when monitoring compliance using SCAP include the following:

- Verifying policy compliance.
- Detecting changes in compliance.

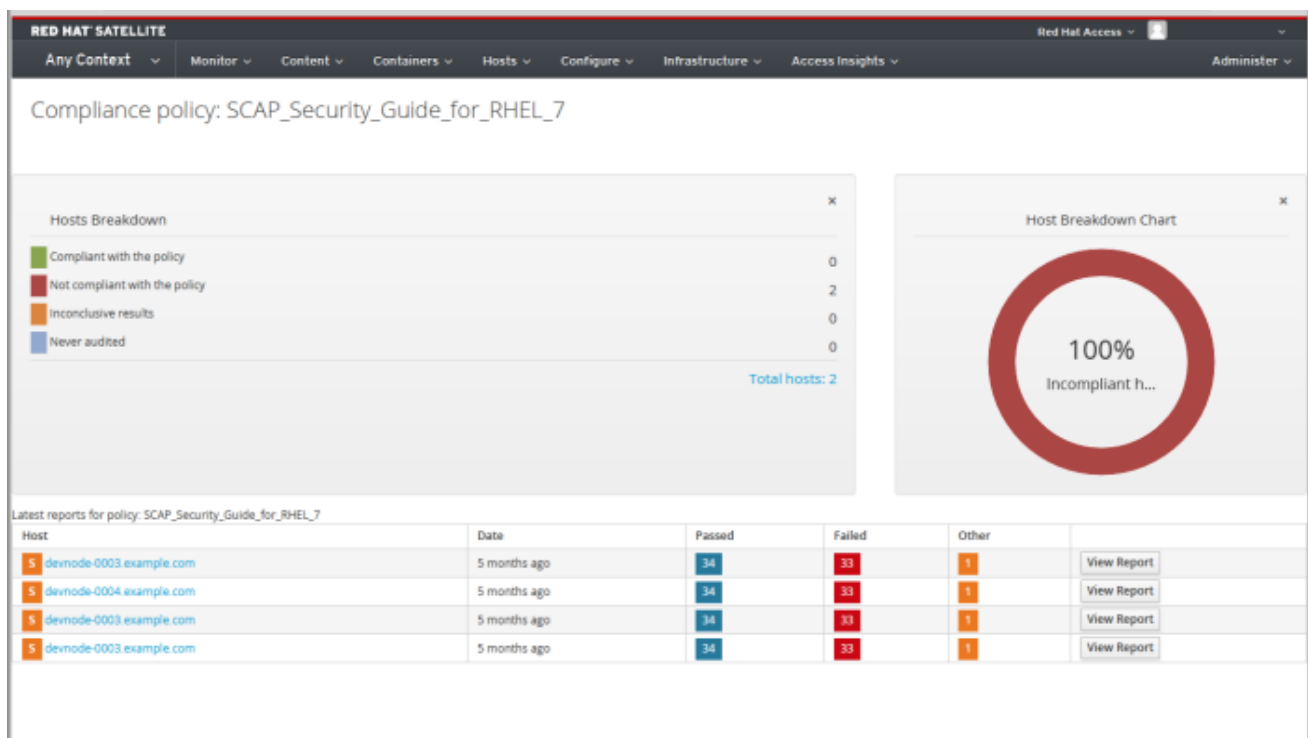
The Satellite web UI provides all the necessary information to achieve these objectives. Verify policy compliance with the compliance policy dashboard. Detect changes in policy compliance by either viewing a compliance report's history or subscribing to notification of changes by email.

6.5.1. Compliance Policy Dashboard

The compliance policy dashboard provides an overview of hosts' compliance with a policy. To view a compliance policy's dashboard, open the Satellite web UI and navigate to **Hosts > Policies**, then click the policy's name. The dashboard provides the following information:

- A ring chart illustrating a high-level view of hosts' compliance with the policy.
- A statistical breakdown of hosts' compliance with the policy, in tabular format.
- Links to the policy's latest report for each host.

The dashboard view provides a statistical summary of hosts' compliance and is a good starting point for compliance management. For all hosts which were evaluated as non-compliant, the **Failed** statistic provides a useful metric for prioritizing compliance effort. Those hosts detected as **Never audited** should also be a priority, since their status is unknown.



6.5.2. Compliance Reports Overview

A compliance report is the output of a policy run against a host. To list all compliance reports, open the Satellite web UI and navigate to **Hosts > Reports**. Each report includes the total number of rules passed or failed per policy. By default, reports are listed in descending date order. To change the sort order, click on the label of the column by which you want it sorted. Click on the same label again to change to either

descending or ascending order. To view an individual report, click **View Report**. To view all reports for a host, or a subset of hosts, use the **Search** field. To delete a compliance report, click the arrow beside **View Report** and select **Delete**.

When managing the policy compliance of hosts, it is useful to monitor compliance changes over time. You can use the **Search** field to narrow the list of reports to one or more hosts and evaluate the changes manually. Alternatively, you can configure notification emails.

6.5.3. Searching Compliance Reports

The Compliance Reports search field allows you to narrow the list of reports. Narrowing your attention on a subset of hosts allows you to focus resources where they are most needed. To apply a filter, enter search criteria in the **Search** field and either press Enter or click **Search**. The search performed is case-insensitive. Click on the empty **Search** field to see a list of available search parameters.

See [Table 5.3, “Supported Operators for Granular Search”](#) for details of all available search operators. You can create complex queries with the logical operators: **and**, **not** and **has**. Regular expressions are not valid search criteria, however multiple fields can be used in a single search expression.

Logical Operators

- **not**: Negates an expression.
- **has**: Object must have a specified property.
- **and**: Combines search criteria.

Search Use Cases

The following search criteria finds all compliance reports for which more than five rules failed.

```
failed > 5
```

The following search criteria finds all compliance reports created after November 5, 2015, for hosts whose host name contains the string **prod-**.

```
host ~ prod- AND date > "Nov 5, 2015"
```

The following search criteria finds all reports generated by the compliance_policy **rhel7_audit** from an hour ago.

```
"1 hour ago" AND compliance_policy = date = "1 hour ago" AND  
compliance_policy = rhel7_audit
```

To again list **all** available compliance reports, delete the **Search** criteria and press **Enter** or click **Search**.

Bookmarking Your Searches

You can bookmark a search, allowing you to apply the same search criteria again.

To Bookmark a Search:

1. Apply your search criteria.
2. From the **Search** list select **Bookmark this search**.

3. Complete the **Name** field.

If you want the bookmark available to other users of this Satellite instance, select the **Public** check box.

4. Click **Submit**.

To use a bookmark, navigate to **Hosts > Reports**, click the drop-down item beside the **Search** button and click the bookmark.

6.5.4. Viewing a Compliance Report

Navigate to **Hosts > Reports** and click **View Report** in the row of the specific host.

A compliance report consists of the following sections:

- Introduction
- Evaluation Characteristics
- Compliance and Scoring
- Rule Overview

6.5.4.1. Evaluation Characteristics

This section provides details about an evaluation against a specific profile, including the host that was evaluated, the profile used in the evaluation, and when the evaluation started and finished. For reference, the IPv4, IPv6, and MAC addresses of the host are also listed.

Evaluation Characteristics

Target machine

The fully-qualified domain name (FQDN) of the evaluated host. Example: **test-system.example.com**.

Benchmark URL

The URL of the SCAP content against which the host was evaluated. Example: **/var/lib/openscap/content/1fbdc87d24db51ca184419a2b6f**.

Benchmark ID

The identifier of the benchmark against which the host was evaluated. A benchmark is a set of profiles. Example: **xccdf_org.ssgproject.content_benchmark_RHEL_7**.

Profile ID

The identifier of the profile against which the host was evaluated. Example: **xccdf_org.ssgproject.content_profile_rht-ccp**.

Started at

The date and time at which the evaluation started, in ISO 8601 format. Example: **2015-09-12T14:40:02**.

Finished at

The date and time at which the evaluation finished, in ISO 8601 format. Example: **2015-09-12T14:40:05**.

Performed by

The local account name under which the evaluation was performed on the host. Example: **root**.

6.5.4.2. Compliance and Scoring

This section provides an overview of whether or not the host is in compliance with the profile's rules, a breakdown of compliance failures by severity, and an overall compliance score as a percentage. If compliance with a rule was not checked, this is categorized in the **Rule results** as **Other**.

6.5.4.3. Rule Overview

This section provides details of every rule and the compliance result, with the rules presented in a hierarchical layout.

Select or clear the check boxes to narrow the list of rules included in the compliance report. For example, if the focus of your review is any non-compliance, clear the **pass** and **informational** check boxes.

To search all rules, enter a criterion in the **Search** field. The search is dynamically applied as you type. The **Search** field only accepts a single plain-text search term and it is applied as a case-insensitive search. When you perform a search, only those rules whose descriptions match the search criterion will be listed. To remove the search filter, delete the search criterion.

For an explanation of each result, hover the cursor over the status shown in the **Result** column.

6.5.4.4. Examining Rule Results

To determine why a host failed compliance on a rule, click on the rule's title. The window which then opens provides further details, including: a description of the rule (with instructions for bringing the host into compliance if available), the rationale for the rule, and in some cases a remediation script.



WARNING

Do not implement any of the recommended remedial actions or scripts without first testing them in a non-production environment.

6.5.5. Compliance Email Notifications

The Satellite Server sends an OpenSCAP Summary email to all users who subscribe to the **Openscap policy summary** email notifications. For more information on subscribing to email notifications see [Section 5.1.6, “Configuring Email Notifications”](#). Each time a policy is run, Satellite checks the results against the previous run, noting any changes between them. The email is sent according to the frequency requested by each subscriber, providing a summary of each policy and its most recent result.

An **OpenSCAP Summary** email message contains the following information:

- Details of the time period it covers.
- Totals for all hosts by status: changed, compliant, and noncompliant.
- A tabular breakdown of each host and the result of its latest policy, including totals of the rules that passed, failed, changed, or where results were unknown.

6.6. SPECIFICATIONS SUPPORTED BY OPENSCAP

The following specifications are supported by OpenSCAP:

Title	Description	Version
XCCDF	The Extensible Configuration Checklist Description Format	1.2
OVAL	Open Vulnerability and Assessment Language	5.11
-	Asset Identification	1.1
ARF	Asset Reporting Format	1.1
CCE	Common Configuration Enumeration	5.0
CPE	Common Platform Enumeration	2.3
CVE	Common Vulnerabilities and Exposures	-
CVSS	Common Vulnerability Scoring System	2.0

CHAPTER 7. BACKING UP AND RESTORING SATELLITE SERVER AND CAPSULE SERVER

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.

7.1. 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 **satellite-backup** script. Backing up to a separate location is recommended. Backing up to a separate storage device on a separate system is highly recommended. Satellite services are unavailable during the backup. The backup can be scheduled for a quiet time using **cron**, see the [Example 7.1, “A Weekly Full Backup Followed by Daily Incremental Backups”](#).

Prerequisites

- 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.
- Ensure that you either encrypt the backups or move them to a secure location to minimize the risk of damage or unauthorized access to hosts.

NOTE

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

```
# katello-service stop
```

Start the services after creating a snapshot or conventional backup:

```
# katello-service start
```

7.1.1. Estimating the size of a Backup

The full backup creates uncompressed archives of MongoDB, PostgreSQL, and Pulp database files and Satellite configuration files. Compression occurs after the archives are created to decrease the time when Satellite services are unavailable. Consequently, a full backup requires space to store the following data:

- Uncompressed Satellite databases and configuration files.
- Compressed Satellite databases and configuration files.
- An extra 20% of the total estimated space to ensure a reliable backup.

To Estimate the Size of a Backup

1. Enter the **du** command to estimate the space required to store uncompressed directories containing Satellite databases and configuration files, for example:

```
# du -csh /var/lib/mongodb /var/lib/pgsql/data /var/lib/pulp \
/etc /root/ssl-build /var/www/html/pub /opt/puppetlabs
480G /var/lib/mongodb
100G /var/lib/pgsql/data
100G /var/lib/pulp
680G total
37M /etc
900K /root/ssl-build
100K /var/www/html/pub
2M /opt/puppetlabs
680GB total
```

In this example, the uncompressed backup data occupies 680 GB in total.



NOTE

The **/opt/puppetlabs** directory is used for Puppet 4. Use **/var/lib/puppet** for Puppet 3.

2. Calculate how much space is required to store the compressed data.
[Table 7.1, “Backup Data Compression Ratio”](#) demonstrates compression ratio of all data items used in the backup.

Table 7.1. Backup Data Compression Ratio

Data type	Directory	Ratio	Example results
MongoDB database files	/var/lib/mongodb	10 - 15%	480 GB → 420 GB
PostgreSQL database files	/var/lib/pgsql/data	15 - 20%	100 GB → 80 GB
Pulp RPM files	/var/lib/pulp	-	(not compressed)
Configuration files	/etc /root-ssl/build /var/www/html/pub /opt/puppetlabs	5 - 10%	50 MB → 45 MB

In this example, the compressed backup data occupies 500 GB in total.

3. Calculate how much space is required to store a backup. Sum up the estimated values of compressed and uncompressed backup data and add extra 20% to ensure a reliable backup. This example requires 680 GB plus 500 GB for the uncompressed and compressed backup data, 1180 GB in total. With 240 GB of extra space, 1420 GB must be allocated for the backup location.

7.1.2. Performing a Full Backup of Satellite Server or Capsule Server

Red Hat Satellite 6.3 uses the **satellite-backup** script to make and restore backups. To see the usage statement, enter a command as follows:

```
# satellite-backup --help
```

From Satellite 6.2.8, the **satellite-backup** creates a time-stamped subdirectory in the backup directory you specify. The **satellite-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 stops and restarts services as required.

To Perform a Full Offline Backup of Satellite Server or Capsule Server:

This procedure performs a full offline backup. Satellite services are unavailable during the backup process.



WARNING

Request other users of Satellite Server or Capsule Server to save any changes and warn them that Satellite services are unavailable for the duration of the backup. Ensure no other tasks are scheduled for the same time as the backup.

1. Ensure your backup location has enough disk space to store the backup. For more information, see [Section 7.1.1, “Estimating the size of a Backup”](#).
2. Run the backup script:

```
# satellite-backup backup_directory
```

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

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

7.1.3. Performing a Backup without Pulp Content

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.

**WARNING**

Request other users of Satellite Server or Capsule Server to save any changes and warn them that Satellite services are unavailable for the duration of the backup. Ensure no other tasks are scheduled for the same time as the backup.

1. Ensure your backup location has enough disk space to store the backup. For more information, see [Section 7.1.1, “Estimating the size of a Backup”](#).
2. Run the backup script:

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

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

7.1.4. Performing an Incremental Backup

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.

**WARNING**

Request other users of Satellite Server or Capsule Server to save any changes and warn them that Satellite services are unavailable for the duration of the backup. Ensure no other tasks are scheduled for the same time as the backup.

1. Ensure your backup location has enough disk space to store the backup. For more information, see [Section 7.1.1, “Estimating the size of a Backup”](#).
2. Run the backup script:
With Pulp content:

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

Without Pulp content:

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

The **satellite-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 creates 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.

7.1.5. Example - A Weekly Full Backup Followed by Daily Incremental Backups

Example 7.1. A Weekly Full Backup Followed by Daily Incremental Backups

An example script that makes a full backup on a Sunday and incremental backups on all other days of the week. Backups are stored in subdirectories named **\$YEAR-\$WEEK**. Each subdirectory contains the full backup from the previous week's Sunday and incremental backups from the subsequent days. This script requires a daily cron job.


```
#!/bin/bash -e
export PATH=/sbin:/bin:/usr/sbin:/usr/bin
DESTINATION=/var/backup
YEAR=$(date +%Y)
WEEK=$(date +%V)
if [[ $(date +%w) == 0 ]]; then
    satellite-backup $DESTINATION/$YEAR-$((WEEK + 1)) --assumeeyes
else
    LAST=$(ls -td -- $DESTINATION/$YEAR-$WEEK/*/ | head -n 1)
    satellite-backup $DESTINATION/$YEAR-$WEEK --incremental "$LAST" --
    assumeeyes
fi
exit 0
```

Note that the **satellite-backup** script requires **/sbin** and **/usr/sbin** directories to be in **PATH**.

7.1.6. Performing an Online Backup

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 repeats 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. This prolongs the backup procedure as the Pulp part of the backup is restarted.



IMPORTANT

Satellite 6 uses two database systems, PostgreSQL and MongoDB. There are records that exist in both PostgreSQL and MongoDB 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.

It is recommended to use a snapshot backup method in production, as described in [Section 7.1.7, “Performing a Snapshot Backup”](#). If you still want to use the online backup method in production, ensure that no modifications occur during the backup.



WARNING

Request other users of Satellite Server or Capsule Server to save any changes and warn them that Satellite services are unavailable for the duration of the backup. Ensure no other tasks are scheduled for the same time as the backup.

1. Ensure your backup location has enough disk space to store the backup. For more information, see [Section 7.1.1, “Estimating the size of a Backup”](#).
2. Run the backup script:

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

7.1.7. Performing a Snapshot Backup

The snapshot backup method uses Logical Volume Manager (LVM) snapshots of the Pulp, MongoDB, and PostgreSQL directories. The actual backup is then created from the LVM snapshots and not from the running Satellite as with online backup, which mitigates the risk of creating an inconsistent backup. The snapshot backup method is faster than a full off-line backup, which reduces Satellite downtime. As such, it suits well for backing up highly populated Satellite servers with long backup times.

To Perform a Snapshot Backup:

This procedure performs a snapshot backup. It can be combined with other **satellite-backup** sub-commands, except an online-backup.

Prerequisites

- The system uses LVM for the snapshotted directories (**/var/lib/pulp/**, **/var/lib/mongodb/**, and **/var/lib/pgsql/**).
- The free disk space in the relevant volume group (VG) is three times the size of the snapshot. More precisely, the VG has to have enough space unreserved by the member logical volumes (LVs) to accommodate new snapshots. In addition, one of the LVs has to have enough free

space for the backup directory.

- The target backup directory is on a different LV than the snapshotted directories.



WARNING

Request other Satellite Server or Capsule Server users to save any changes and warn them that Satellite services are unavailable for the duration of the backup. Ensure no other tasks are scheduled for the same time as the backup.

Run the backup script:

```
# satellite-backup --snapshot backup_directory
```

The **satellite-backup** script stops all services which could impact the backup. After the successful backup, all services are restarted and LVM snapshots are removed.

7.2. 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 7.1, “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 is 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 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.
- Ensure that you run the **satellite-restore** script as **root**.
- Ensure that all SELinux contexts are correct. Enter the following command to restore the correct SELinux contexts:

```
# restorecon -Rnv /
```

To Restore Satellite Server or Capsule Server from a Full Backup:

1. Install Satellite 6 using the procedures in [Installing Satellite Server](#) in the *Installation Guide*.
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:

```
# satellite-restore backup_directory
```

Where *backup_directory* is the time-stamped directory or subdirectory containing the backed-up data. The target directory is read from the configuration files contained within the archive. If the target directory does not exist when trying to recover, it gives an error and asks 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 [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.

To Restore Satellite Server or Capsule Server from an Incremental Backup:

1. Install Satellite 6 using the procedures in [Installing Satellite Server](#) in the *Installation Guide*.
2. Restore the last full backup as described in [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:

```
# satellite-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 is read from the configuration files contained within the archive. If the target directory does not exist when trying to recover, it gives an error and asks for the correct directory.

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

7.3. BACKING UP AND RESTORING CAPSULE SERVER USING A VIRTUAL MACHINE SNAPSHOT

There are three methods of backing up Capsule Server:

- Using the **satellite-backup** script as described in [Section 7.1, “Backing up Satellite Server or Capsule Server”](#). The **satellite-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 of a virtual machine with your Capsule Server on it as described below. Note that this method is different from the snapshot backup method described in [Section 7.1.7, “Performing a Snapshot Backup”](#).

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 **satellite-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 *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.

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**.

7.4. RENAMING A SATELLITE SERVER OR CAPSULE SERVER

Renaming a Satellite Server or Capsule Server requires use of the **satellite-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.

7.4.1. 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.

If you use external authentication, you must reconfigure Satellite Server for external authentication after you run the **satellite-change-hostname** script. The **satellite-change-hostname** script breaks external authentication for Satellite Server. For more information about configuring external authentication, see [Chapter 10, Configuring External Authentication](#)

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 *Red Hat Satellite Installation Guide*.
- Backup the Satellite Server. The **satellite-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 [Section 10.2, "Using Identity Management"](#).

Rename a Satellite Server

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

```
# satellite-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 *Red Hat Satellite 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 more information, see [Configuring and Running Remote Commands](#) in *Managing Hosts*.

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 more information, see [Configuring and Running Remote Commands](#) in *Managing Hosts*.

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

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

6. On the Satellite Server, synchronize content for each Capsule Server.

- a. List all Capsule Servers with their ID numbers:

```
# hammer capsule list
```

- b. Enter the following command for each Capsule Server:

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

7.4.2. 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 *Red Hat Satellite Installation Guide*.
- Backup the Capsule Server. The **satellite-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 7, Backing Up and Restoring Satellite Server and Capsule Server](#).

Renaming a Capsule Server:

1. On 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 /root/new-capsule.example.com-certs.tar
```

Ensure that you enter the full path to the **.tar** file.

- 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 *Red Hat Satellite Installation Guide*.

2. On 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 **satellite-change-hostname** script, providing the host's new name, Satellite credentials, and certificates archive filename.

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

Ensure that you enter the full path to the **.tar** file.

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 *Red Hat Satellite 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.

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

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

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

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

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

```
# subscription-manager refresh
```

```
# yum repolist
```

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.

CHAPTER 8. MAINTAINING 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.

8.1. LOGGING AND REPORTING

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

Table 8.1. 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](#)^[1] for more information.

8.2. ENABLING DEBUG LOGGING

This section describes how to enable *debug logging* to provide detailed debugging information for Satellite 6.3. Debug logging provides the most detailed log information and can help with troubleshooting issues that can arise with Satellite 6.3 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
```

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
```

8.3. COLLECTING INFORMATION FROM LOG FILES

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

Table 8.2. 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.

8.4. 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 8.8.5, “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*

8.5. CLEANING UNUSED TASKS

Cleaning unused tasks reduces disk space in the database and limits the rate of disk growth. When you perform regular cleaning, Satellite backup completes faster and overall performance is higher.

To clean unused tasks

The installer has a feature to enable a cron job to automatically remove old tasks. This feature is not enabled by default to avoid unwanted tasks cleanup.

1. Enable the cron job:

```
# satellite-installer --foreman-plugin-tasks-automatic-cleanup true
```

2. By default, the cron job is scheduled to run every day at 19:45. To adjust to the time, change the value of the **--foreman-plugin-tasks-cron-line** parameter:

```
# satellite-installer --foreman-plugin-tasks-cron-line "00 15 * * *"
```

The previous command schedules the cron job to run every day at 15:00, see **man 5 crontab** for more details on cron format.

To change the period after which to delete all the tasks and to configure further advanced settings of the cron job, change the content of the **/etc/foreman/plugins/foreman-tasks.yaml** file.

8.6. RECOVERING FROM A FULL DISK

The following procedure describes how to resolve the situation when a logical volume (LV) with the Pulp database on it has no free space.

To recover from a full disk

1. Let running Pulp tasks finish but do not trigger any new ones as they can fail due to the full disk.
2. Ensure that the LV with the **/var/lib/pulp** directory on it has sufficient free space. Here are some ways to achieve that:
 - a. Remove orphaned content:

```
# foreman-rake katello:delete_orphaned_content  
RAILS_ENV=production
```

This is run weekly so it will not free much space.

- b. Change the download policy from **Immediate** to **On Demand** for as many repositories as possible and remove already downloaded packages. See the Red Hat Knowledgebase

solution [How to change syncing policy for Repositories on Satellite 6.2 from "Immediate" to "On-Demand"](#) on the Red Hat Customer Portal for instructions.

- c. Grow the file system on the LV with the `/var/lib/pulp` directory on it. For more information, see [Growing a File System on a Logical Volume](#) in the *Red Hat Enterprise Linux 7 Logical Volume Manager Administration Guide*.



NOTE

If you use an untypical file system (other than for example ext3, ext4, or xfs), you might need to unmount the file system so that it is not in use. In that case:

- Stop Katello services:

```
# katello-service stop
```

- Grow the file system on the LV.

- Start Katello services:

```
# katello-service start
```

3. If some Pulp tasks failed due to the full disk, run them again.

8.7. RECLAIMING DISK SPACE FROM MONGODB

The MongoDB database can use a large amount of disk space especially in heavily loaded deployments. The following procedure describes how to reclaim some of this disk space.

Prerequisites

- A backup of the MongoDB database. For instructions on creating a backup, see [Section 7.1.3, "Performing a Backup without Pulp Content"](#).
- Pulp services are stopped:

```
# systemctl stop goferd httpd pulp_workers pulp_celerybeat \
pulp_resource_manager pulp_streamer
```

To reclaim disk space from MongoDB

1. Access the MongoDB shell:

```
# mongo pulp_database
```

2. Check the amount of disk space used by MongoDB before a repair:

```
> db.stats()
```

3. Ensure that you have free disk space equal to the size of your current MongoDB database plus 2 GB. If the volume containing the MongoDB database lacks sufficient space, you can mount a separate volume and use that for the repair.

4. Enter the repair command:

```
> db.repairDatabase()
```

Note that the repair command blocks all other operations and can take a long time to complete, depending on the size of the database.

5. Check the amount of disk space used by MongoDB after a repair:

```
> db.stats()
```

6. Start Pulp services:

```
# systemctl start goferd httpd pulp_workers pulp_celerybeat \
pulp_resource_manager pulp_streamer
```

8.8. 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.

8.8.1. 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.

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.

8.8.2. 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.

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 8.8.5, “Creating Support Cases Using the Red Hat Access Plug-in”](#).

8.8.3. 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.

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:
 - In the **Search** field, provide a key word or phrase.
 - From the drop-down list, choose a specific **Case Group**. Your organization has defined **Case Groups** inside the Red Hat Customer Portal.
 - 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.

8.8.4. 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.

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

1. Complete the instructions from [Section 8.8.3, “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**.

8.8.5. 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.


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.

8.9. 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.

8.10. 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.

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

CHAPTER 9. MONITORING CAPSULE SERVER

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

9.1. 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. For more information, see [Installing Capsule Server](#) in the *Red Hat Satellite Installation Guide*.

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.

9.2. 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 9.3, “Monitoring Puppet”](#)).

9.3. 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 10. 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 10.1, “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 10.2, “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 10.3.1, “Using Active Directory with Cross-Forest Trust”](#).
- Using direct AD as an external identity provider. For more information see [Section 10.3.2, “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 automatically manages the life cycle of any system registered to a realm or domain provider. See [Section 10.5, “External Authentication for Provisioned Hosts”](#) for more information.

10.1. USING LDAP

10.1.1. 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 [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 6.3](#) 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
```

10.1.2. 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 [Using Identity Management](#) or [Using Active Directory](#) for more information on those options.

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_yppbind 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 10.1.3, "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 10.1.3, "LDAP Setting Descriptions and Examples"](#) for examples.
7. On the **Locations** tab, select locations from the left table. Selected locations are 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 are 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 [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, the Satellite administrator(s) have to assign roles manually. See [Assigning Roles to a User](#) to assign user accounts appropriate roles in Satellite.

10.1.3. LDAP Setting Descriptions and Examples

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

Table 10.1. 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>
Account password	<p>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.</p>
Base DN	<p>The top level domain name of the LDAP directory.</p>

Setting	Description
Groups base DN	The top level domain name of the LDAP directory tree that contains groups.
LDAP filter	A filter to restrict LDAP queries. See Section 10.1.3.1, “Example LDAP Filters” for examples.
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 may see a Permissions Denied warning. These users have to contact their Satellite administrator to request that suitable roles are associated with their accounts.
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 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 10.2. 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
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 10.3. 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 10.4. 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
First name attribute	givenName
Surname attribute	sn
Email address attribute	mail

10.1.3.1. Example LDAP Filters

As an administrator, you can create LDAP filters to restrict the access of specific users to Satellite.

Table 10.5. Example filters for allowing specific users to login

User	Filter
User1, User3	(memberOf=cn=Group1,cn=Users,dc=domain,dc=example)

User	Filter
User2, User3	(memberOf=cn=Group2,cn=Users,dc=domain,dc=example)
User1, User2, User3	(&(objectClass=user)((memberOf=cn=Group1,cn=Users,dc=domain,dc=example)(memberOf=cn=Group2,cn=Users,dc=domain,dc=example)))

LDAP directory structure

The LDAP directory structure that the filters in the example use:

```

DC=Domain, DC=Example
|
|----- CN=Users
|
|----- CN=Group1
|----- CN=Group2
|----- CN=User1
|----- CN=User2
|----- CN=User3

```

LDAP group membership

The group membership that the filters in the example use:

Group	Members
Group1	User1, User3
Group2	User2, User3

10.2. USING IDENTITY MANAGEMENT

Select from one of the following methods:

- [Section 10.2.1, “Using Identity Management Directly”](#)
- [Section 10.2.2, “Using Identity Management with LDAP Authentication”](#)

10.2.1. 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](#)^[2].

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

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](#)^[3].

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](#)^[4].

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](#)^[5].

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 ^[6].
5. Configure the IdM integration with the Satellite Server as described in [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
```

10.2.2. Using Identity Management with LDAP Authentication

To attach Identity Management as an external authentication source with no single sign-on support, see [Section 10.1, “Using LDAP”](#) for more information.

10.3. USING ACTIVE DIRECTORY

Select from one of the following methods:

- [Section 10.3.1, “Using Active Directory with Cross-Forest Trust”](#)
- [Section 10.3.2, “Using Active Directory Directly”](#)
- [Section 10.3.3, “Using Active Directory with LDAP Authentication”](#)

10.3.1. 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](#)^[7].

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 10.2, “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
```

10.3.2. Using Active Directory Directly

This section shows how to use direct Active Directory (AD) as an external authentication source for Satellite Server. Direct AD integration means that 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 [To Enroll Satellite Server with the AD Server](#), then finalize the AD integration with use of GSS-proxy as described in [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 Red Hat Satellite Server to be deployed on Red Hat Enterprise Linux 7.1 or later.

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
```

To Enroll Satellite Server with the AD Server:

1. Install the required packages:

```
# yum install sssd adcli realmd ipa-python-compat krb5-workstation
```

2. Enroll 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 a previously installed Satellite or during 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.

To Configure Direct AD Integration with GSS-proxy:

1. Create the **/etc/ipa/** directory and the **default.conf** file:

```
# mkdir /etc/ipa
# touch /etc/ipa/default.conf
```

2. To the **default.conf** file, add the following content:

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

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

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

4. Create the **/etc/gssproxy/00-http.conf** file with the following content:

```
[service/HTTP]
mechs = krb5
cred_store = keytab:/etc/krb5.keytab
cred_store = ccache:/var/lib/gssproxy/clients/krb5cc_%U
euid = 48
```

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

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

6. Create a keytab entry:

```
# KRB5_KTNAME=FILE:/etc/httpd/conf/http.keytab net ads keytab add
HTTP -U administrator -d3 -s /etc/net-keytab.conf
# chown root.apache /etc/httpd/conf/http.keytab
# chmod 640 /etc/httpd/conf/http.keytab
```

7. Enable IPA authentication in Satellite:

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

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

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

9. Configure the Apache server to use the gssproxy service:

- a. Create the **/etc/systemd/system/httpd.service** file with the following content:

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

- b. Apply changes to the service:

```
# systemctl daemon-reload
```

10. 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.

To confirm that SSO is working as expected, on Satellite Server, enter the following command to retrieve the Kerberos ticket of the LDAP user:

```
# kinit ldapuser
```

To view the Kerberos ticket, enter the following command:

```
# klist
```

To view output from successful SSO-based authentication, enter the following command:

```
# curl -k -u : --negotiate
https://satellite.example.com/users/extlogin
```

```
<html><body>You are being <a
href="https://satellite.example.com/users/4-ldapuserexample-
com/edit">redirected</a>.</body></html>
```

Users can now configure Kerberos SSO in their browsers to be able to log in without filling in access credentials in the Satellite UI. For more information on configuring the Firefox browser see the [Red Hat Enterprise Linux System-Level Authentication Guide](#). If you use the Internet Explorer browser, add 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](#)^[8].

10.3.3. Using Active Directory with LDAP Authentication

To attach Active Directory as an external authentication source with no single sign-on support, see [Section 10.1, “Using LDAP”](#) for more information. For an example configuration, see [How to configure Active Directory authentication with TLS on Satellite 6](#).

10.4. 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 10.1, “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 10, 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.

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 would 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.

10.5. 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.

10.5.1. 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.

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` prepares 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 require your Identity Management server configuration details for 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 external Capsule Server, use **satellite-installer --scenario capsule** to configure the realm:

```
# satellite-installer --scenario capsule \
--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
```

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:
 - On the **Realm** subtab, provide the realm name, the type of realm to use and the realm proxy.
 - On the **Locations** subtab, choose the locations where the new realm is intended for use.
 - 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 details the steps on how to automatically add hosts to an IdM host group.

10.5.2. 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.

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:

- *hostgroup_name* is the host group's name.
- *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:

- **automember-add** flags the group as an automember group.
- **--type=hostgroup** identifies that the target group is a host group, not a user group.
- *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:

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

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

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

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

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

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

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

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

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

CHAPTER 11. 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.

Plug-ins for Satellite typically include the word **foreman** in the RPM package name and plug-ins for Capsule include **smart_proxy** in the name. View the RPM package description to confirm the identity of a plug-in using **yum info** or **rpm -qi**. For information on upstream plug-ins, see the [Plugins](#) section of the *Foreman* website.

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.

11.1. ADDING ADDITIONAL PLUG-INS

To list the plug-ins available from the configured repositories, you can search using part of the package name. For example, on Satellite Server, 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 on a Satellite, enter as **root**:

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

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

```
# yum list installed | grep proxy
```

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

```
# yum install tfm-rubygem-foreman_templates
```

Restart Katello services for the plug-in to be registered:

```
# katello-service restart
```

For additional information on plug-ins, see the [Popular Plugins](#) and [List of Plugins](#) sections on the *Foreman* website.



IMPORTANT

Support is 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 <http://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/el_X_/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 Katello services for the plug-in to be registered:

```
# katello-service restart
```

11.2. 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](#)^[9].

- On Red Hat Enterprise Linux 7, see [Fixing Problems](#)^[10].

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__l
ogger.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
/usr/share/foreman/config/hooks/host/managed/before_provision/_10__l
ogger.sh
```

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

```
# chmod u+x
/usr/share/foreman/config/hooks/host/managed/before_provision/_10__l
ogger.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 Katello services for the hook to be registered:

```
# katello-service restart
```

Every Foreman or Rail object can have a hook. See the **/usr/share/foreman/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

11.2.1. 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 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

11.2.2. 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](#)^[11] documentation.

11.2.3. 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 is 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 is sufficient for most users. Otherwise, closing standard input is recommend 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 is used as the priority value, thereby influencing when it is executed in relation to DNS, DHCP, VM creation, and other tasks.

11.2.4. Hook Failures and Rollback

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

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

[10] 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

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

APPENDIX A. SETTINGS PARAMETERS

This section further describes some parameters in the **Administer > Settings** page.

Provisioning tab:

Type of name generator

Specifies the method used to generate a host name when creating a new host. The default **Random-based** option generates an unique random host name which you can but do not have to use. This is useful for users who create many hosts and do not know how to name them. The **MAC-based** option is for bare-metal hosts only. If you delete a host and create it later on, it receives the same host name based on the MAC address. This can be useful for users who recycle servers and want them to always get the same host name. The **Off** option disables the name generator function and leaves the host name field blank.

Safemode rendering

Enables safe mode rendering of provisioning templates. The default and recommended option **Yes** denies the access to variables and any object that is not whitelisted within Satellite. When set to **No**, any object may be accessed by a user with permission to use templating features, either via editing of templates, parameters or smart variables. This permits users full remote code execution on Satellite Server, effectively disabling all authorization. This is not a safe option, especially in bigger companies.

General tab:

Fix DB cache

Satellite maintains a cache of permissions and roles. When set to **Yes**, Satellite recreates this cache on the next restart. The default option is **No**.

Use Gravatar

When set to **Yes**, Satellite displays user avatars by matching the user email address with an email address at <https://gravatar.com>. The default and for security reasons recommended option is **No**.