



Red Hat CloudForms 4.5

General Configuration

A guide to configuring and tuning Red Hat CloudForms

Red Hat CloudForms 4.5 General Configuration

A guide to configuring and tuning Red Hat CloudForms

Red Hat CloudForms Documentation Team
cloudforms-docs@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 configuring CloudForms Management Engine, including appliance settings, access control, web console appearance, and registration. Information and procedures in this book are relevant to CloudForms Management Engine administrators. This guide is for the Red Hat CloudForms 4.5 and Red Hat CloudForms Management Engine version 5.8. If you have a suggestion for improving this guide or have found an error, please submit a Bugzilla report at <http://bugzilla.redhat.com> against Red Hat CloudForms Management Engine for the Documentation component. Please provide specific details, such as the section number, guide name, and CloudForms version so we can easily locate the content.

Table of Contents

CHAPTER 1. SETTINGS OVERVIEW	6
CHAPTER 2. MY SETTINGS	7
2.1. VISUAL SETTINGS	7
2.1.1. Grid and Tile Icons	7
2.1.2. Setting Default Items Per Page	8
2.1.3. Setting the Start Page	8
2.1.4. Setting Display Settings	9
2.2. DEFAULT VIEWS	9
2.2.1. Setting Default Views for the User Interface	10
2.2.2. Setting Default Views for Services	10
2.2.3. Setting Default Views for Clouds	11
2.2.4. Setting Default Views for Infrastructure Components	12
2.2.5. Setting Default Views for Containers	13
2.3. DEFAULT FILTERS	14
2.3.1. Setting Default Filters for Cloud	14
2.3.2. Setting Default Filters for Containers	15
2.3.3. Setting Default Filters for Infrastructure	15
2.3.4. Setting Default Filters for Services	15
2.4. TIME PROFILES	15
2.4.1. Creating a Time Profile	15
2.4.2. Editing a Time Profile	17
2.4.3. Copying a Time Profile	17
2.4.4. Deleting a Time Profile	17
CHAPTER 3. TASKS	19
3.1. MY VM AND CONTAINER ANALYSIS TASKS	19
3.1.1. Filtering the VM and Container Analysis Task List	19
3.1.2. Deleting a Task from the VM and Container Analysis Task List	19
3.2. MY OTHER UI TASKS	20
3.2.1. Filtering the My Other UI Tasks List	20
3.2.2. Deleting a My Other UI Task	20
3.3. ALL VM AND CONTAINER ANALYSIS TASKS	20
3.3.1. Filtering the VM and Container Analysis Task List	21
3.3.2. Deleting a VM and Container Analysis Task	21
3.4. ALL OTHER TASKS	21
3.4.1. Filtering the All Other Tasks List	22
3.4.2. Deleting a Task from the All Other Tasks List	22
CHAPTER 4. CONFIGURATION	23
4.1. SETTINGS	23
4.1.1. Regions	24
4.1.1.1. Region Scope	27
4.1.1.2. Region Settings	27
4.1.1.3. Capacity and Utilization Collections	27
4.1.1.3.1. Capacity and Utilization Collection Settings	27
4.1.1.3.2. Enabling a Cluster, Host, or Datastore for Capacity and Utilization Collection	28
4.1.1.4. Tags	28
4.1.1.4.1. Company Tag Categories and Tags	28
4.1.1.4.2. Creating a Tag Category	28
4.1.1.4.3. Deleting a Tag Category	30
4.1.1.4.4. Creating a Company Tag	30

4.1.1.4.5. Deleting a Company Tag	30
4.1.1.4.6. Importing Tags for Virtual Machines	31
4.1.1.4.7. Importing Tags for a Virtual Machine from a CSV File	31
4.1.1.4.8. Importing Custom Values for Virtual Machines and Hosts	31
4.1.1.4.9. Importing Asset Tags for a Virtual Machine from a CSV File	32
4.1.1.5. Registering Red Hat CloudForms	32
4.1.1.5.1. Registering Appliances	33
4.1.1.5.2. Subscription Management for Virtual Environments	34
4.1.2. Profiles	35
4.1.2.1. Creating an Analysis Profile	35
4.1.2.2. Creating a Host Analysis Profile	35
4.1.2.3. Creating a Virtual Machine Analysis Profile	36
4.1.2.4. Editing an Analysis Profile	38
4.1.2.5. Copying an Analysis Profile	38
4.1.2.6. Setting a Default Analysis Profile	38
4.1.3. Zones	39
4.1.3.1. Creating a Zone	39
4.1.3.2. Deleting a Zone	40
4.1.3.3. Editing a Zone	41
4.1.3.4. Adding SmartProxy Affinity to a Zone:	41
4.1.4. Servers	41
4.1.4.1. Changing Server Settings	42
4.1.4.1.1. Basic Information Settings	42
4.1.4.1.2. Server Control Settings	43
4.1.4.1.3. Server Roles	43
4.1.4.1.4. VMware Console Settings	48
4.1.4.1.5. NTP Servers Settings	49
4.1.4.1.6. Configuring SNMP	49
4.1.4.1.7. Outgoing SMTP Email Settings	50
4.1.4.1.8. Web Services Settings	51
4.1.4.1.9. Logging Settings	52
4.1.4.1.10. Custom Support URL Settings	52
4.1.4.2. Authentication	52
4.1.4.2.1. Changing Authentication Settings	52
4.1.4.2.2. Configuring LDAP Authentication	53
4.1.4.2.3. Trusted Forests	56
4.1.4.2.4. Assigning CloudForms Account Roles Using LDAP Groups	57
4.1.4.2.5. Using Existing LDAP Groups to Assign Account Roles	57
4.1.4.2.6. Using Groups Named by CloudForms to Assign Account Roles	59
4.1.4.2.7. Configuring Lookup Priority for LDAP Groups	60
4.1.4.2.8. Testing LDAP Configuration	60
4.1.4.2.9. Amazon Settings	61
4.1.4.2.10. External Authentication (httpd)	61
4.1.4.2.11. Configuring External Authentication Using SAML	64
4.1.4.3. Workers	70
4.1.4.3.1. Changing Settings for a Worker	70
4.1.4.4. Database	70
4.1.4.4.1. Changing a Database Setting	70
4.1.4.5. Customization and Logos	71
4.1.4.5.1. Custom Logos	71
4.1.4.5.2. Uploading a Custom Logo to the User Interface	71
4.1.4.5.3. Customizing the Login Background	72
4.1.4.5.4. Customizing the Login Panel Text	73

4.1.4.5.5. Displaying the Custom Configuration Settings	73
4.1.4.6. Advanced Settings	74
4.1.4.6.1. Editing Configuration Files Manually	74
4.1.4.6.2. Configuration Parameters	74
4.1.5. Schedules	96
4.1.5.1. Scheduling SmartState Analyses and Backups	96
4.1.5.1.1. Scheduling a SmartState Analysis or Compliance Check	97
4.1.5.2. Scheduling a Database Backup	98
4.1.5.2.1. Modifying a Schedule	100
4.2. ACCESS CONTROL	100
4.2.1. Creating a Tenant	100
4.2.2. Creating a Project	100
4.2.3. Managing Tenant and Project Quotas	101
4.2.4. Tagging Tenants and Projects	101
4.2.5. Creating a User	102
4.2.6. Deleting a User	102
4.2.7. Groups	103
4.2.8. Creating a Group	103
4.2.9. Roles	104
4.2.9.1. Account Roles and Descriptions	105
4.2.10. Creating a Role	107
4.3. DIAGNOSTICS	108
4.3.1. Region Diagnostics	108
4.3.1.1. Server Role Priorities	109
4.3.2. Region Aware Server Roles	109
4.3.3. Setting the Priority of a Failover Role	110
4.3.4. Zone Diagnostics	110
4.3.4.1. Viewing the Status of Server Roles	110
4.3.4.2. Zone Aware Server Roles	110
4.3.4.2.1. Removing an Inactive Server	111
4.3.4.3. Zone Log Collections	111
4.3.4.3.1. Setting the Location of the Log Depot	112
4.3.4.3.2. Collecting and Downloading Logs from All Servers in a Zone	113
4.3.4.4. Capacity and Utilization Repair	113
4.3.4.4.1. Repairing Capacity and Utilization Data	113
4.3.5. Server Diagnostics	114
4.3.5.1. Workers	114
4.3.5.1.1. Reloading Worker Display	114
4.3.5.1.2. Restarting a Worker	115
4.3.5.2. Server and Audit Logs	115
4.3.5.2.1. Collecting Server Logs and Configuration Files	115
4.3.5.2.2. Setting the Location of the Log Depot for a Specific Server	115
4.3.5.2.3. Collecting the Current Log Set of a Server	116
4.3.5.2.4. Collecting All Log Sets from a Server	116
4.3.5.2.5. Viewing the Server, Audit, and Production Logs	117
4.3.5.2.6. Viewing the Server Log	117
4.3.5.2.7. Reloading the Server Log	117
4.3.5.2.8. Downloading the Server Log	117
4.3.5.2.9. Viewing the Audit Log	118
4.3.5.2.10. Reloading the Audit Log	118
4.3.5.2.11. Downloading the Audit Log	118
4.3.5.2.12. Viewing the Production Log	118
4.3.5.2.13. Reloading the Production Log	119

4.3.5.2.14. Downloading the Production Log	119
4.4. DATABASE OPERATIONS	119
4.4.1. Viewing Information on the VMDB	119
4.4.2. Database Regions and Replication	120
4.4.2.1. Creating a Region	120
4.4.3. Configuring Database Replication and Centralized Administration	121
4.4.3.1. Configuring a Remote Copy	121
4.4.3.2. Configuring the Global Copy	122
4.4.3.3. Resetting Database Replication	123
4.4.4. Backing Up and Restoring a Database	123
4.4.4.1. Running a Single Database Backup	123
4.4.4.2. Restoring a Database from a Backup	124
4.4.5. Performing a Binary Backup and Restoring the Database	125
4.4.5.1. Performing a Binary Backup	125
4.4.5.2. Restoring a Database from the Backup	126
4.4.6. Running Database Garbage Collection	127
4.4.7. Changing the Database Password	127
4.4.7.1. Changing the Password on the Database Appliance	127
4.4.7.2. Changing the Password on the Worker Appliances	128
4.4.8. Adding a New Appliance to an Existing Region with a Non-default Password	129
4.4.9. Configuring Scheduled Database Maintenance	129
CHAPTER 5. SMARTPROXIES	131
5.1. INSTALLING THE SMARTPROXY FROM THE CONSOLE	131
5.2. ENTERING CREDENTIALS AND OPERATING SYSTEM FOR THE TARGET HOST	131
CHAPTER 6. ABOUT	133
CHAPTER 7. RED HAT INSIGHTS	134
7.1. OVERVIEW TAB	134
7.2. RULES TAB	135
7.2.1. States	135
7.2.2. Info Listed	135
7.3. SYSTEMS TAB	136

CHAPTER 1. SETTINGS OVERVIEW

The settings menu is located at the top right of the Red Hat CloudForms user interface.

To view or modify global settings for your appliance, click your username to open the settings dropdown menu, and click **My Settings** to modify. The availability of each menu item depends on the role assigned to your user account. For more information on roles, see [Section 4.2.9, “Roles”](#).

The following is a list of the menu items available from the settings menu:

- **My Settings**

This menu is available to all Red Hat CloudForms users. The settings in this menu control how elements in the user interface are displayed, time profiles, and tags for the currently logged-in individual user.

- **Tasks**

This menu allows you to view virtual machine SmartState Analysis tasks that can be tracked through the console. The status of each task is displayed, including time started, time ended, what part of the task is currently running, and any errors encountered.

- **Configuration**

This menu allows you to specify enterprise, region, zone, and server settings for your Red Hat CloudForms infrastructure. Diagnostics such as logs and the status of processes are also shown here. The **Configuration** menu is available only to super administrators and administrators.

CHAPTER 2. MY SETTINGS

The options under the **My Settings** menu allow you to configure options specific to the user account with which you log in to the Red Hat CloudForms user interface, such as the default view displayed on login, and personal tags. You can also configure the color scheme, button options, and external RSS feeds on the main dashboard.

2.1. VISUAL SETTINGS

The options under the **Visual** menu allow you to configure how user interface elements are displayed in the web user interface. For all options, click **Save** to update the settings, or click **Reset** to undo any unsaved changes that have been made on the current screen.

2.1.1. Grid and Tile Icons

This group of settings is used to control the view of your virtual thumbnails. Each thumbnail can be viewed as a single icon or as an icon with four quadrants. Use the quadrant view to see a component's properties at a glance.

Use the following procedure to change grid and tile icons:

1. From the settings menu, navigate to **My Settings**, then click on the **Visual** tab.
2. In **Grid/Tile Icons**, set items to **ON** to display all four quadrants for the item, or **OFF** to display only one icon.
3. Click **Save**.

Grid/Tile Icons

Show Infrastructure Provider Quadrants

ON

Show Cloud Provider Quadrants

ON

Show Host Quadrants

ON

Show Datastores Quadrants

ON

Show VM Quadrants

ON

Show Template Quadrants

ON

Truncate Long Text

Middle (AB...34) ▾

- Set **Show Infrastructure Provider Quadrants** to **ON** to see the four icons in your provider under **Compute** → **Infrastructure** → **Providers**.

- Set **Show Cloud Provider Quadrants** to **ON** to see the four icons in your hosts under **Cloud → Providers**.
- Set **Show Host Quadrants** to **ON** to see the four icons in your hosts under **Compute → Infrastructure → Hosts**.
- Set **Show Datastores Quadrants** to **ON** to see the four icons in your datastores under **Compute → Infrastructure → Datastores**.
- Set **Show VM Quadrants** to **ON** to see the four icons in your virtual machines under **Compute → Infrastructure → Virtual Machines**.
- Set **Show Template Quadrants** to **ON** to see the four icons in your templates under **Compute → Infrastructure → Virtual Machines → Templates**.

Set any of the above options to **OFF** to see only one icon instead of four quadrants.

- Use the **Truncate Long Text** list to specify how the names of items are displayed if they are too long to show in full. Select the option based on the pattern shown.

2.1.2. Setting Default Items Per Page

Use the following procedure to set the default number of items to display on each resource page.

1. From the settings menu, navigate to **My Settings**, then click on the **Visual** tab.
2. In the **Default Items Per Page** area, select the default number of items to display for each view from the corresponding drop down list.

Default Items Per Page

Grid View	20	▼
Tile View	20	▼
List View	20	▼
Reports	20	▼

3. Click **Save**.

2.1.3. Setting the Start Page

Use the following procedure to set the default start page after logging in. For example, instead of going to the **Red Hat CloudForms** dashboard, you can set the default start page to see a list of your virtual machines.

1. From the settings menu, navigate to **My Settings**, then click on the **Visual** tab.

2. In the **Start Page** area, select the page to display at login.

Start Page

Show at Login

Cloud Intelligence / Dashboard ▾

3. Click **Save**.

2.1.4. Setting Display Settings

Use the following procedure to set your own themes, colors, and time zone for the console. These settings are specific to the logged-on user.

1. From the settings menu, navigate to **My Settings**, then click on the **Visual** tab.
2. Make selections from **Display Settings** for the following items:

Display Settings

Chart Theme

MIQ ▾

Time Zone

(GMT+00:00) UTC ▾

Locale

Global Default ▾

- a. Use **Chart Theme** to select a group of colors and font sizes specifically for charts.
- b. Use **Time Zone** to select the time zone in which to display the console.



NOTE

In time zones where clocks are set forward for daylight savings time, the time zone correctly displays as EDT (Eastern Daylight Time) in the console. When the clocks are set back, it correctly displays as EST (Eastern Standard Time).

- c. Use **Locale** to select the language in which to display the console.
3. Click **Save**.

2.2. DEFAULT VIEWS







The options under the **Default View** menu allow you to configure the default layout used to display individual screens in the Red Hat CloudForms user interface. The options you select under this menu specify the default options for each screen, but you can also change the layout for each screen using the layout buttons on each screen.

2.2.1. Setting Default Views for the User Interface

Use the following procedure to set general view options:

1. From the settings menu, navigate to **My Settings**, then click on the **Default Views** tab.
2. In the **General** area, click the appropriate button for the way you want to view each type of screen listed. The selected view shows as a blue icon.



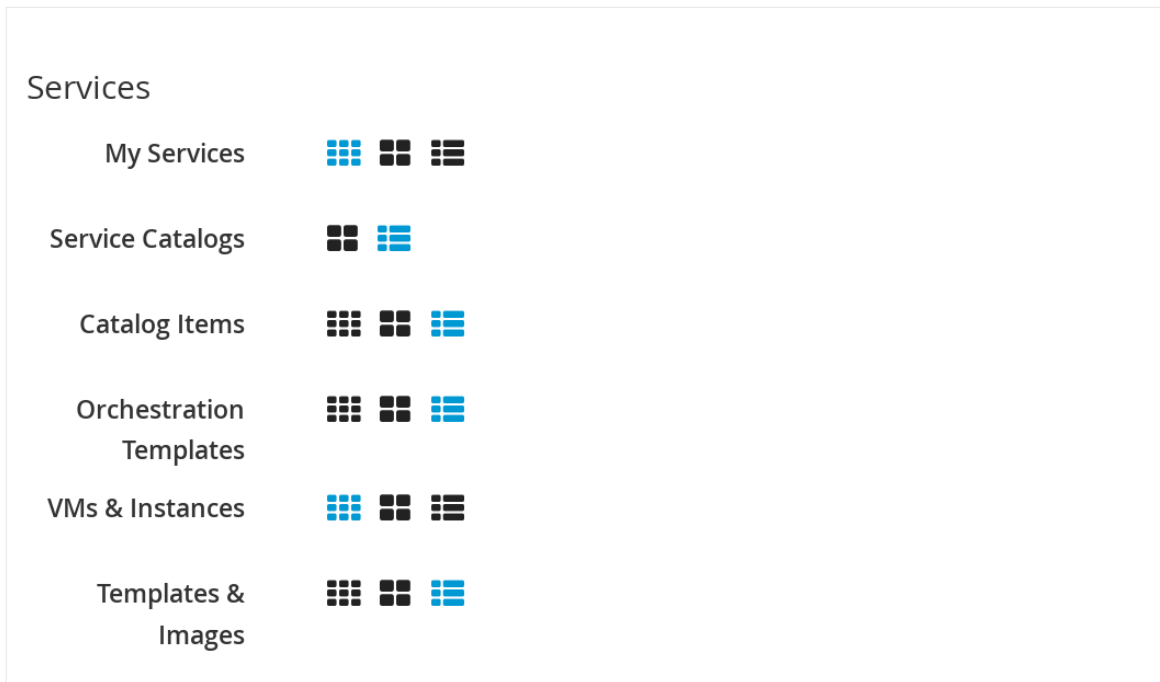
3. Click  (**Grid View**) to view virtual thumbnails or icons.
4. Click  (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
5. Click  (**List View**) to view a detailed text listing.
6. Click  (**Expanded View**) for an expanded view.
7. Click  (**Compressed View**) for a compressed view.
8. Click  (**Exists Mode**) to view only whether an attribute exists or not.
9. Click **Save**.

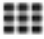


2.2.2. Setting Default Views for Services

Use the following procedure to set default views for services in the **Compute → Services** tab.

1. From the settings menu, navigate to **My Settings**, then click on the **Default Views** tab.

2. In the **Services** area, click the appropriate button for the way you want to view each item.



- Click  (**Grid View**) to view virtual thumbnails or icons.
- Click  (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
- Click  (**List View**) to view a text listing.

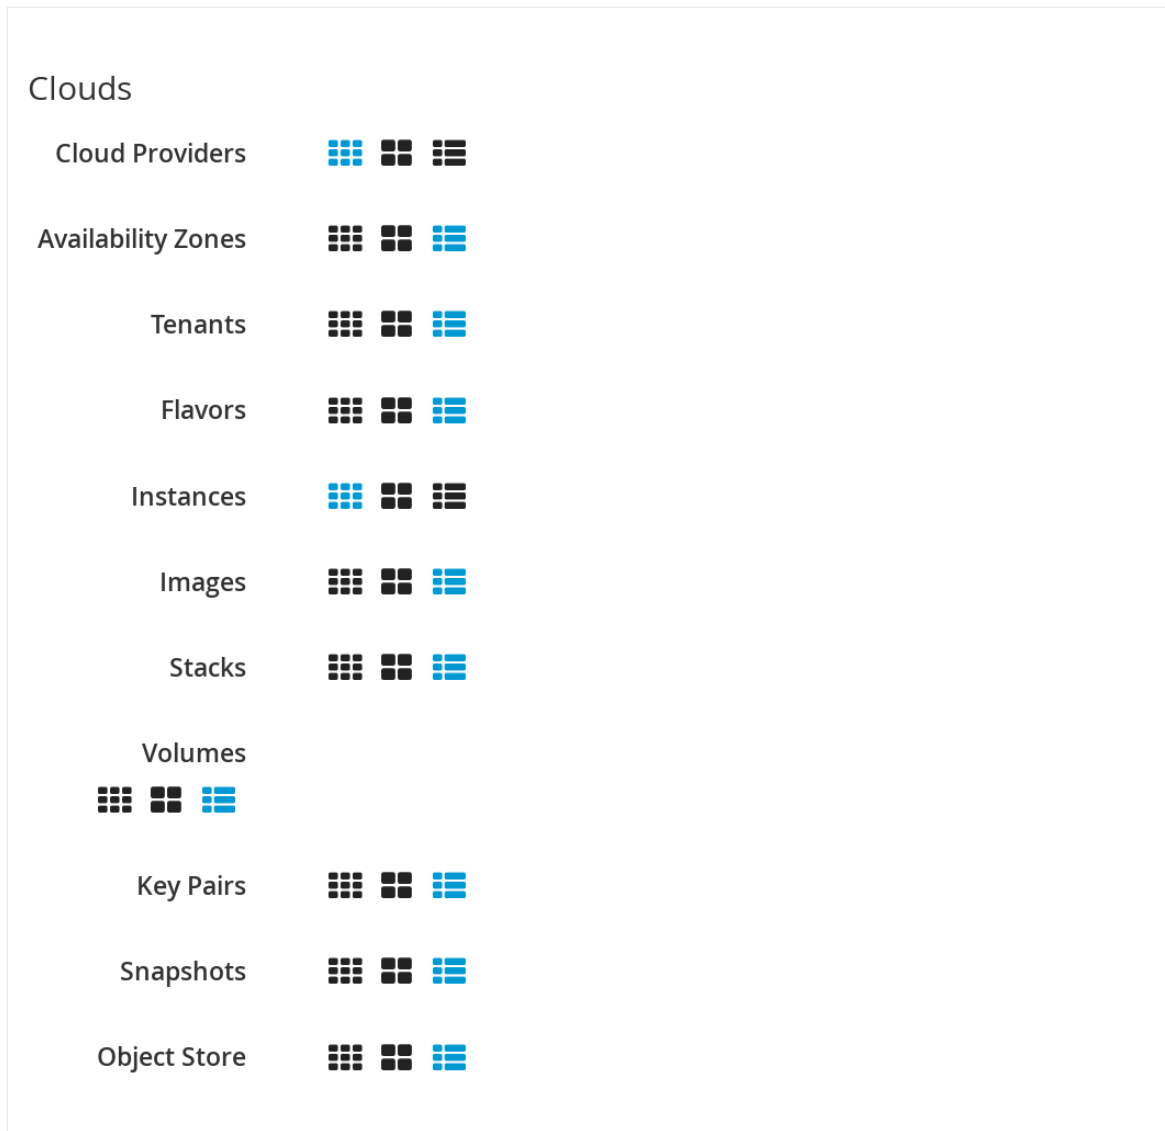
3. Click **Save**.




2.2.3. Setting Default Views for Clouds

Use the following procedure to set default views for clouds in the **Compute → Clouds** tab.

1. From the settings menu, navigate to **My Settings**, then click on the **Default Views** tab.

2. In the **Clouds** area, click the appropriate button for the way you want to view each item.



- Click  (**Grid View**) to view virtual thumbnails or icons.
- Click  (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
- Click  (**List View**) to view a detailed text listing.

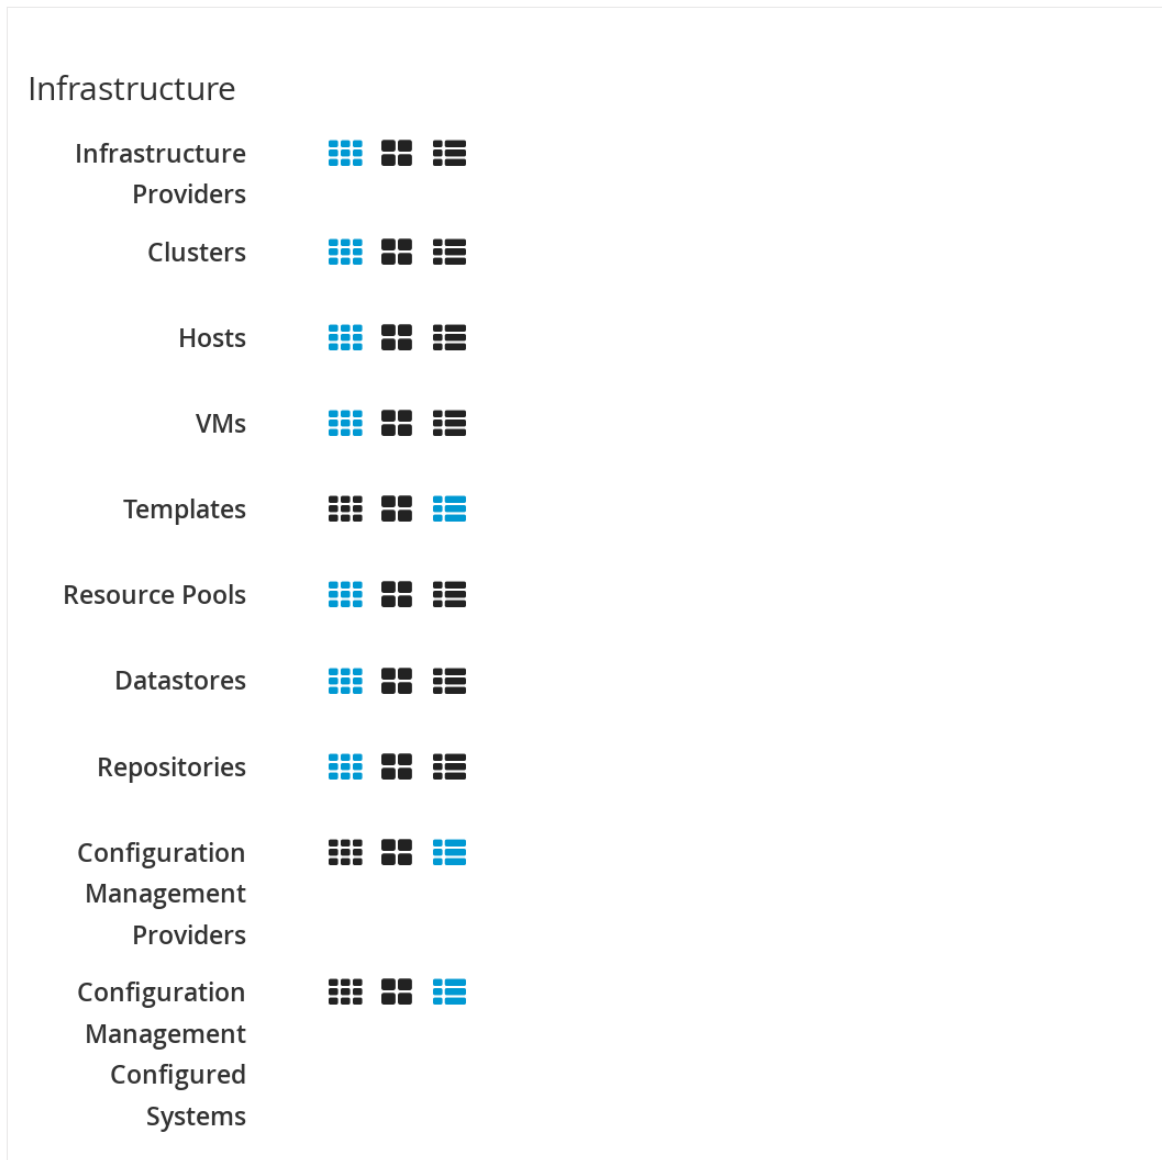
3. Click **Save**.

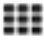


2.2.4. Setting Default Views for Infrastructure Components

Use the following procedure to set default views for infrastructure components in the **Compute** → **Infrastructure** tab.

1. From the settings menu, navigate to **My Settings**, then click on the **Default Views** tab.

2. In the **Infrastructure** area, click the appropriate button for the way you want to view each item.



- Click  (**Grid View**) to view virtual thumbnails or icons.
- Click  (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
- Click  (**List View**) to view a detailed text listing.

3. Click **Save**.

2.2.5. Setting Default Views for Containers

Use the following procedure to set default views for containers in the **Compute** → **Containers** tab.

1. From the settings menu, navigate to **My Settings**, then click on the **Default Views** tab.

2. In the **Containers** area, click the appropriate button for the way you want to view each item.



- Click (**Grid View**) to view virtual thumbnails or icons.
- Click (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
- Click (**List View**) to view a text listing.

3. Click **Save**.

2.3. DEFAULT FILTERS

The options the **Default Filters** menu allow you to configure the default filters displayed for your hosts, virtual machines, and templates. These settings are available to all users.

2.3.1. Setting Default Filters for Cloud

To set default filters for cloud components:

1. From the settings menu, navigate to **My Settings**, then click on the **Default Filters** tab.
2. From the **Cloud** folder, check the boxes for the default filters that you want available. Items that have changed show in blue text.
3. Click **Save**.

2.3.2. Setting Default Filters for Containers

To set default filters for containers:

1. From the settings menu, navigate to **My Settings**, then click on the **Default Filters** tab.
2. From the **Containers** folder, check the boxes for the default filters that you want available. Items that have changed show in blue text.
3. Click **Save**.

2.3.3. Setting Default Filters for Infrastructure

To set default filters for infrastructure components:

1. From the settings menu, navigate to **My Settings**, then click on the **Default Filters** tab.
2. In the **Infrastructure** folder, select the default filters that you want available. Items that have changed show in blue text.
3. Click **Save**.

2.3.4. Setting Default Filters for Services

To Set Default Filters for Services:



1. From the settings menu, navigate to **My Settings**, then click on the **Default Filters** tab.
2. In the **Services** folder, select the default filters that you want available. Items that have changed show in blue text.
3. Click **Save**.

2.4. TIME PROFILES

The options under the **Time profiles** menu allow you to specify the hours for which data is displayed when viewing capacity and utilization screens. Time profiles are also used to configure performance and trend reports, and for **Optimize** pages.

2.4.1. Creating a Time Profile

To create a time profile:

1. From the settings menu, navigate to **My Settings**, then click on the **Time Profiles** tab.
2. Click  (**Configuration**), and  (**Add a new Time Profile**).

Time Profile Information

Description	Peak Work South America												
Scope	Current User ▼												
Days	(All)	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday					
	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>					
Hours	(All)	AM: 12-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
		PM: 12-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12
		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Timezone	(GMT-11:00) International Date L ▼												
Roll Up Daily Performance	<input type="checkbox"/>												

3. Type a meaningful name in the **Description** field.
4. Select the users who can access the time profile from the **Scope** list:
 - Select **All Users** to create a time profile that is available to all users. Only the super administration and administration roles can create, edit, and delete a global profile.
 - Select **Current User** if this time profile should only be available to the user creating it.
5. Check the **Days** and **Hours** for the time profile.
6. For **Timezone**, you can select a specific time zone or, you can let the user select a time zone when displaying data.
7. If you select a specific time zone, you also have the option to **Roll Up Daily Performance** data. This option is only available to users with the administration or super administration role. Enabling the **Roll Up Daily Performance option** reduces the time required to process daily capacity and utilization reports and to display daily capacity and utilization charts.
8. Click **Add**.



NOTE



The following relationships exist between time zones and performance reports:

- The configured time zone in a performance report is used to select rolled up performance data, regardless of the user's selected time zone.
- If the configured time zone is null, it defaults to UTC time for performance reports.
- If there is no time profile with the report's configured time zone that is also set to roll up capacity and utilization data, the report does not find any records.

For non-performance reports, the user's time zone is used when displaying dates and times in report rows.



2.4.2. Editing a Time Profile

To edit a time profile:

1. From the settings menu, navigate to **My Settings**, then click on the **Time Profiles** tab.
2. Check the time profile you want to edit.
3. Click  (**Configuration**), and  (**Edit selected Time Profile**).
4. Make the required changes.
5. Click **Save**.

2.4.3. Copying a Time Profile



To copy a time profile:

1. From the settings menu, navigate to **My Settings**, then click on the **Time Profiles** tab.
2. Check the time profile you want to copy.
3. Click  (**Configuration**), and  (**Copy selected Time Profile**).
4. Make the required changes.
5. Click **Save**.

2.4.4. Deleting a Time Profile

To delete a time profile:

1. From the settings menu, navigate to **My Settings**, then click on the **Time Profiles** tab.

2. Check the time profile you want to delete.
3. Click  (**Configuration**), and  (**Delete selected Time Profiles**).
4. Click **Save**.

CHAPTER 3. TASKS

The options under the **Tasks** menu allow you to view and control currently running tasks in Red Hat CloudForms. The status of each task is displayed, including time started, time ended, what part of the task is currently running, and any errors encountered.

3.1. MY VM AND CONTAINER ANALYSIS TASKS

The **My VM and Container Analysis Tasks** menu allows you to view all tasks running on virtual machines and containers for the currently logged-in user.

From the **My VM and Container Analysis Tasks** menu, you can:

- See jobs that the logged on user created for the SmartProxy either through a schedule or by manually initiating a SmartState Analysis of a virtual machine or container.
- See if a job completed successfully, resulted in an error, or is running.
- See the reason for an error.
- Filter the tasks by status and state.
- View the owner or host of the virtual machine or container referenced.
- Delete a task either explicitly or older than another task.

3.1.1. Filtering the VM and Container Analysis Task List

This procedure describes how to filter the virtual machine and container analysis task list. You can filter the task list by zone, time period, task status, and task state.

To filter the virtual machine and container analysis task list:

1. From the settings menu, select **Tasks**.
2. Click **My VM and Container Analysis Tasks**.
3. From the **Zone** list, select a specific zone, or select **<All Zones>**.
4. From the **24 Hour Time Period** list, select the period of time to view the tasks.
5. For **Task Status**, select the check boxes next to the status to view.
6. From the **Tasks State** list, select the state to view.
7. Click **Apply**.

3.1.2. Deleting a Task from the VM and Container Analysis Task List

To delete a task from the list of virtual machine and container analysis tasks:

1. From the settings menu, select **Tasks**.
2. Click **My VM and Container Analysis Tasks**.

3. Select the tasks to delete. From the **Delete** menu:
 - a. Click **Delete** to delete the selected tasks.
 - b. Click **Delete Older** to delete the tasks older than the selected tasks.
 - c. Click **Delete All** to delete all tasks.
4. Click **OK**.

3.2. MY OTHER UI TASKS

The **My Other UI Tasks** menu allows you to view all tasks running in the user interface for the currently logged-in user.

3.2.1. Filtering the My Other UI Tasks List

This procedure describes how to filter the user interface analysis task list. You can filter the task list by time period, task status, and task state.

To filter the other user interface task list:

1. From the settings menu, select **Tasks**.
2. Click **My Other UI Tasks**.
3. From the **24 Hour Time Period** list, select the period of time to view the tasks.
4. For **Task Status**, select the check boxes next to the status to view.
5. From the **Tasks State** list, select the state to view.
6. Click **Apply**.

3.2.2. Deleting a My Other UI Task

To delete a task from the list of other user interface tasks:

1. From the settings menu, select **Tasks**.
2. Click **My Other UI Tasks**.
3. Select the tasks to delete. From the **Delete** menu:
 - a. Click **Delete** to delete the selected tasks.
 - b. Click **Delete Older** to delete the tasks older than the selected tasks.
 - c. Click **Delete All** to delete all tasks.
4. Click **OK**.

3.3. ALL VM AND CONTAINER ANALYSIS TASKS

The **All VM and Container Analysis Tasks** menu allows you to view all tasks running on virtual machines and containers for all users.

From the **All VM and Container Analysis Tasks** menu, you can:

- See jobs that all users have created for the SmartProxy either through a schedule or by manually initiating a SmartState Analysis of a virtual machine or container.
- See if a job completed successfully, resulted in an error, or is running.
- See the reason for an error.
- Filter the tasks by status and state.
- View the owner or host of the virtual machine or container referenced.
- Delete a task either explicitly or older than another task.

3.3.1. Filtering the VM and Container Analysis Task List

This procedure describes how to filter virtual machine analysis task lists. You can filter the task list by zone, time period, task status, and task state.

To filter the virtual machine and container analysis task list:

1. From the settings menu, select **Tasks**.
2. Click **All VM and Container Analysis Tasks**.
3. From the **Zone** list, select a specific zone, or select **<All Zones>**.
4. From the **24 Hour Time Period** list, select the period of time to view the tasks.
5. For **Task Status**, select the check boxes next to the status to view.
6. From the **Tasks State** list, select the state to view.
7. Click **Apply**.

3.3.2. Deleting a VM and Container Analysis Task

To delete a task from the list of all virtual machine and container analysis tasks:

1. From the settings menu, select **Tasks**.
2. Click **All VM and Container Analysis Tasks**.
3. Select the tasks to delete. From the **Delete** menu:
 - a. Click **Delete** to delete the selected tasks.
 - b. Click **Delete Older** to delete the tasks older than the selected tasks.
 - c. Click **Delete All** to delete all tasks.
4. Click **OK**.

3.4. ALL OTHER TASKS

The **All Other Tasks** menu allows you to view all tasks for all users. This menu is only accessible when you are logged in as the **admin** user.

3.4.1. Filtering the All Other Tasks List

This procedure describes how to filter the list of all other tasks. You can filter the task list by time period, task status, and task state.

To filter the all other tasks list:

1. From the settings menu, select **Tasks**.
2. Click **All Other Tasks**.
3. From the **24 Hour Time Period** list, select the period of time to view the tasks.
4. For **Task Status**, select the check boxes next to the status to view.
5. From the **Tasks State** list, select the state to view.
6. Click **Apply**.

3.4.2. Deleting a Task from the All Other Tasks List

To delete a task from the list of all other tasks:

1. From the settings menu, select **Tasks**.
2. Click **All Other Tasks**.
3. Select the tasks to delete. From the **Delete** menu:
 - a. Click **Delete** to delete the selected tasks.
 - b. Click **Delete Older** to delete the tasks older than the selected tasks.
 - c. Click **Delete All** to delete all tasks.
4. Click **OK**.

CHAPTER 4. CONFIGURATION

The options under the **Configuration** menu allow you to configure global options for your Red Hat CloudForms environment, view diagnostic information, and view analytics on the servers in the environment. The menu displays the Red Hat CloudForms environment at the enterprise, zone, and server levels.

There are four main areas:


- **Settings**
This menu allows you to configure global settings for your Red Hat CloudForms infrastructure. You can also create analysis profiles and schedules for these profiles.
- **Access Control**
This menu contains options for configuring users, groups, roles, and tenants.
- **Diagnostics**
This menu displays the status of your servers and their roles and provides access to logs.
- **Database**
specify the location of your Virtual Machine Database (VMDB) and its login credentials.







4.1. SETTINGS


The options under the **Settings** area provide a hierarchical view of options that allow you to configure global options for the infrastructure of your Red Hat CloudForms environment. At the top level, you have **Settings** including users, LDAP Groups, account roles, capacity and utilization collection, tag categories, values, and imports, custom variable imports, and license uploads. When you click on **Settings** and expand it, you can configure **Analysis Profiles**, **Zones**, and **Schedules**.


My Settings
Tasks
Configuration
About

▼ Settings

▼  CFME Region: Region 1 [1]

- ▼  Analysis Profiles
 -  host default
 -  host sample
 -  sample
- ▼  Zones
 - ▼  Zone: Default Zone (current)

 Server: EVM [1000000000001] (current)

 Schedules

When you go the **Settings** accordion, you are automatically taken to the server list under **Zones**.

4.1.1. Regions

Use **Regions** for centralizing data which is collected from public and private virtualization environments. A region is ultimately represented as a single database for the VMDB. Regions are particularly useful when multiple geographical locations need to be managed as they enable all the data collection to happen at each particular location and avoid data collection traffic across slow links between networks.

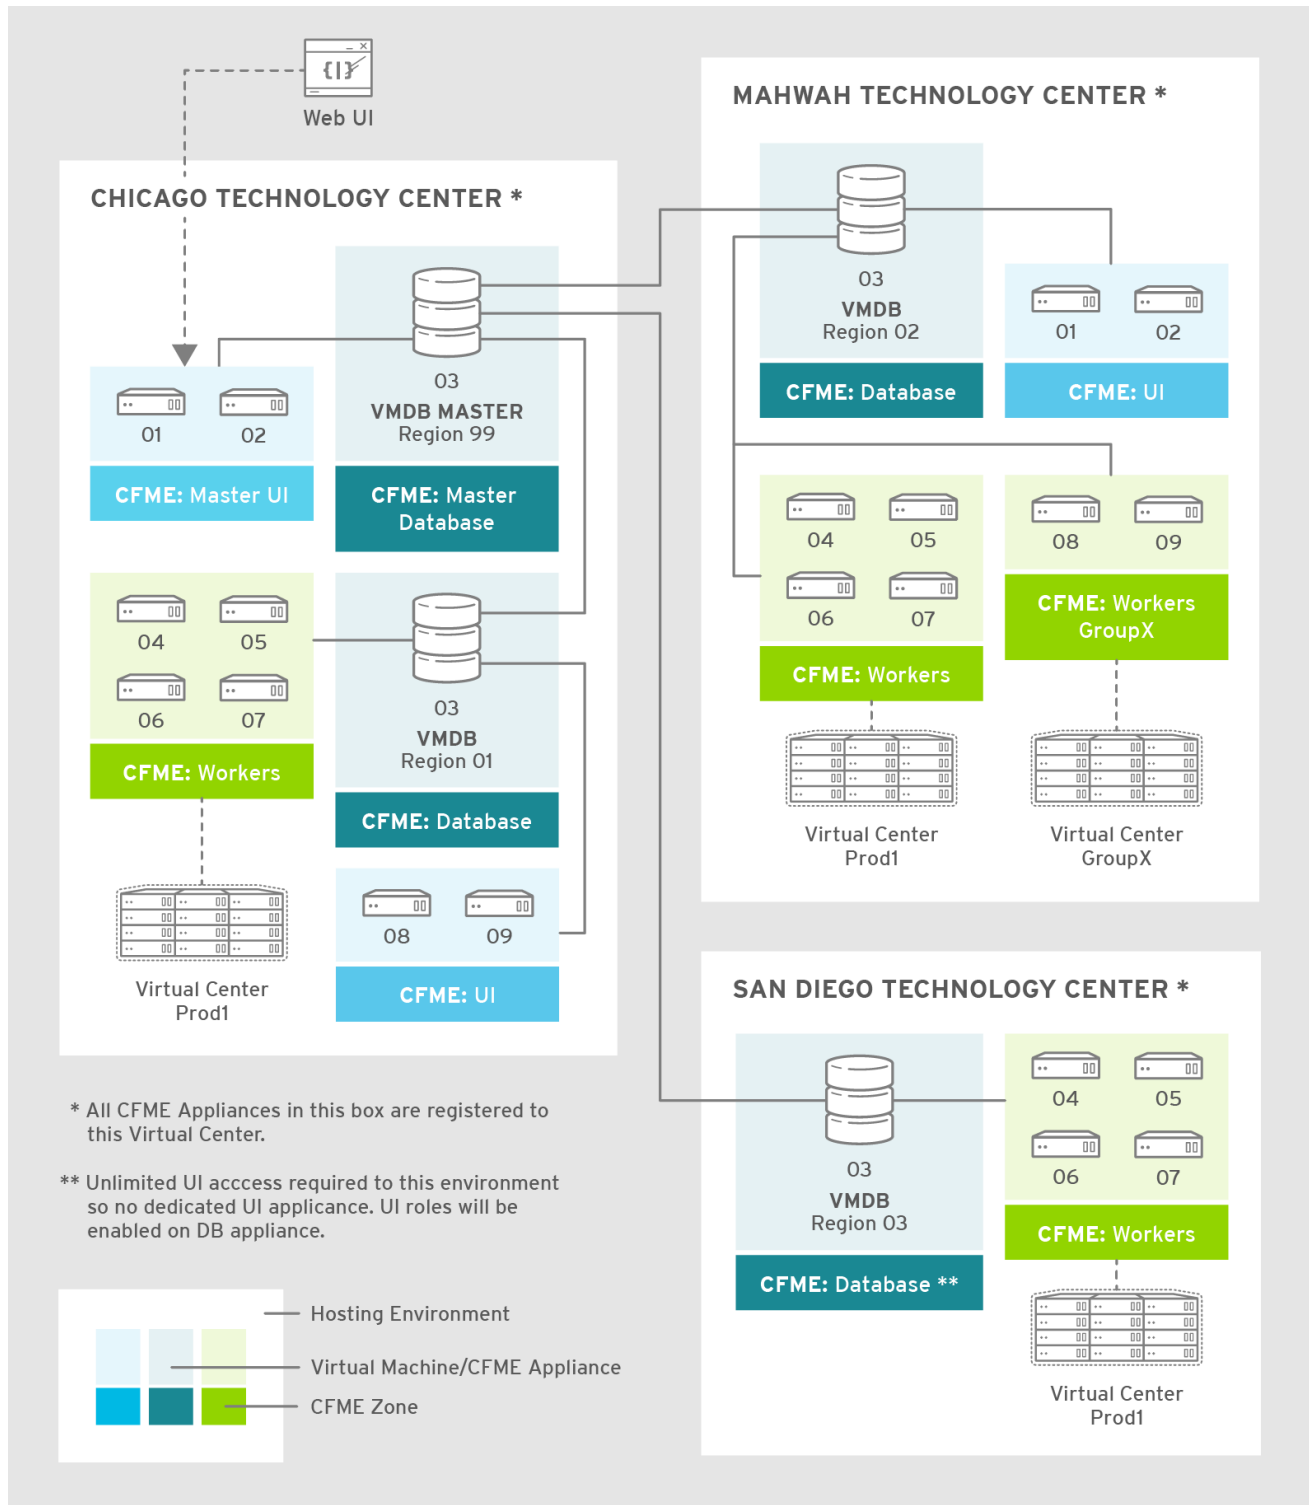
When multiple regions are being used, each with their own unique ID, a master region can be created to centralize the data of all the children regions into a single master database. To do this, configure each child region to replicate its data to the master region database (Red Hat recommends use of region 99). This parent and child region is a one-to-many relationship.

Regions can contain multiple zones, which in turn contain appliances. Zones are used for further segregating network traffic along with enabling failover configurations. Each appliance has the capability to be configured for a number of specialized server roles. These roles are limited to the zone containing the appliance they run on.

Only one failover type of each server role can run in a zone. If multiple appliances have the same failover role, the extras are used as backups that activate only if the primary appliance fails. Non-failover server roles can run on multiple appliances simultaneously in a

zone, so resources can be adjusted according to the workload those roles are responsible for.

The following diagram demonstrates an example of the multiple regions working together in a Red Hat CloudForms environment.

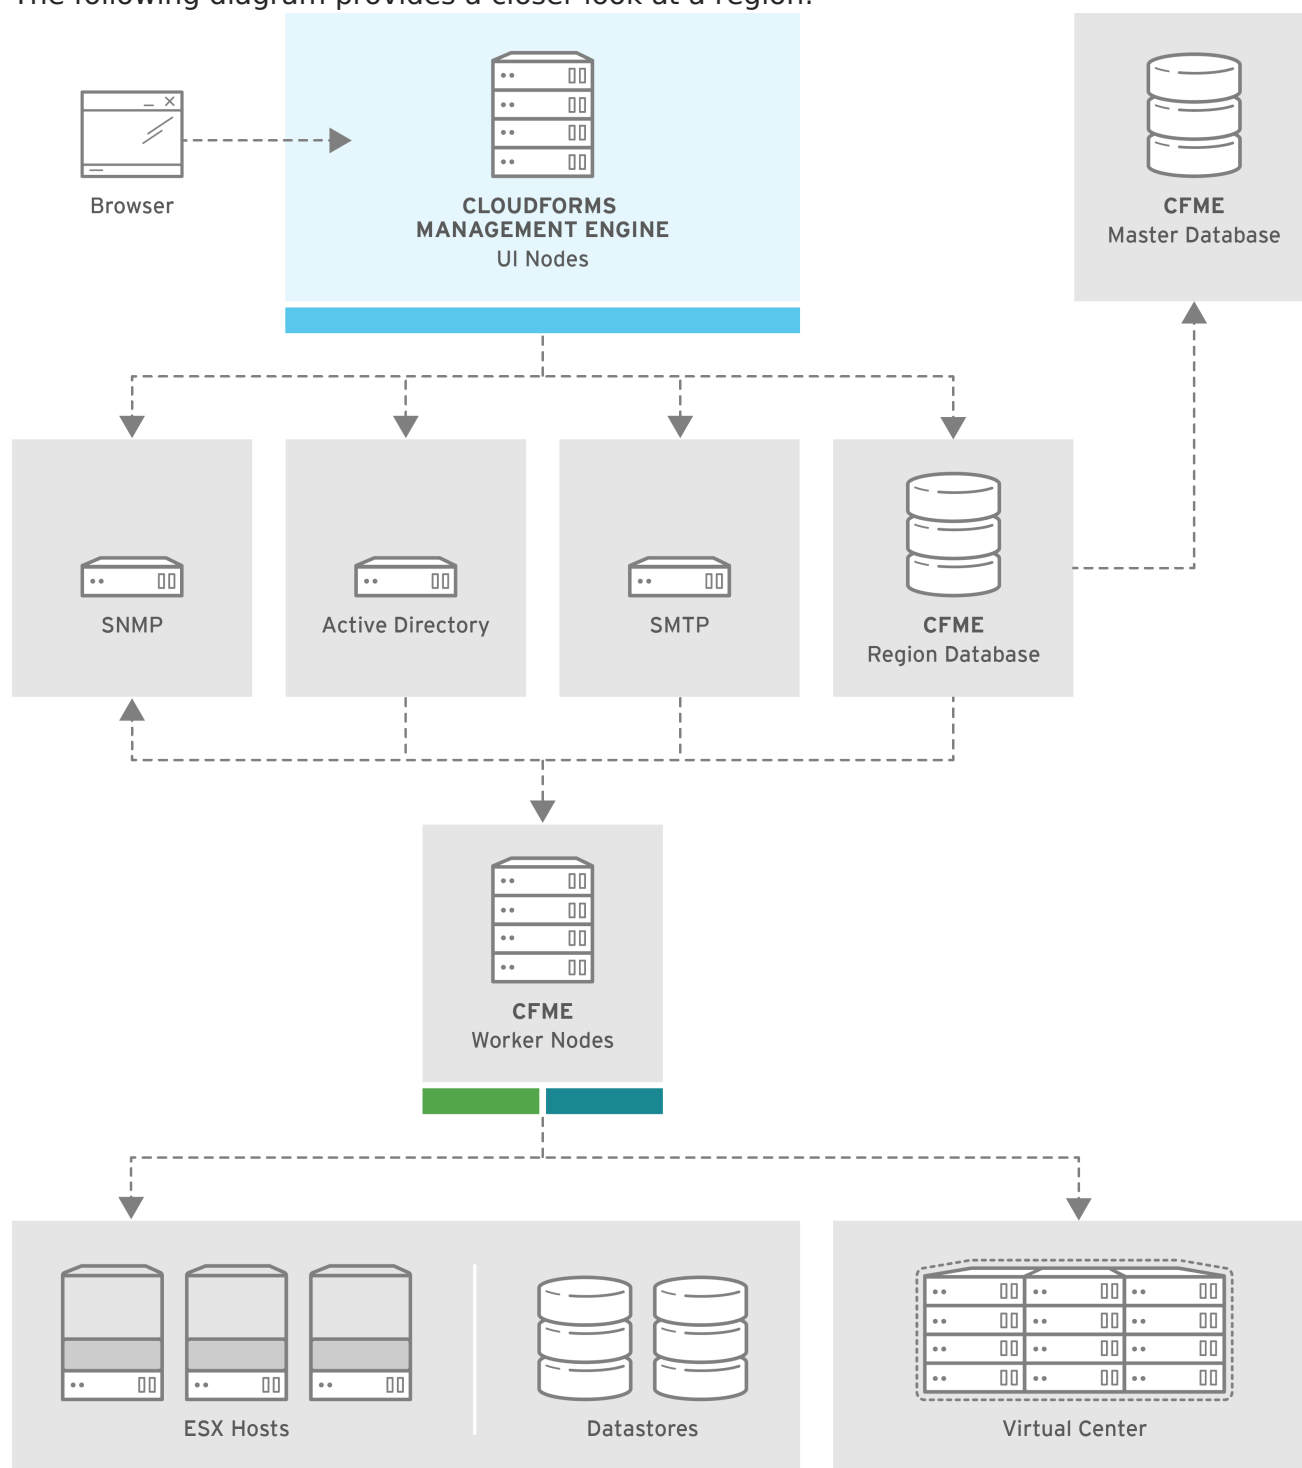


The Master appliance is located in Chicago and contains a master region and a subregion that manages the worker appliances. The Mahwah technology center contains a single subregion that manages two zones. Likewise the San Diego technology center contains a single subregion managing a single zone.

**NOTE**

- Replicating a parent region to a higher-level parent is not supported.
- Parent regions can be configured after the child regions are online.

The following diagram provides a closer look at a region:



CFME_337199_0215

In this region, we have several Red Hat CloudForms appliances acting as UI nodes and worker nodes. These worker nodes execute tasks on the providers in your environment. The Region also uses a region database that reports to a master database on the main Red Hat

CloudForms appliance. All appliances can connect to the authentication services (Active Directory, LDAP, Identity Management), outgoing mail (SMTP), and network services (SNMP).

4.1.1.1. Region Scope

Regions are used to consolidate data from multiple VMDBs to a central database. The database at the top level, the master VMDB, cannot be used for operational tasks such as SmartState Analysis or Capacity and Utilization data collection. It is intended for use as a reporting database that includes all information across multiple subordinate regions. The subordinate regions replicate their information to the master.



NOTE

The subordinate regions are not aware of each other from a database perspective. You cannot see information from one subordinate region in another. The only VMDB with data visibility to all subordinate regions is the top level.

Master Regions Scope

- Reports all information from all subordinate VMDBs reporting up to it.
- Can perform power operations on virtual machines from subordinate regions.
- Controls its own access control list.

Subordinate Regions Scope

- Each subordinate controls its own access control independent of the other regions.
- Can only do work (such as SmartState Analysis and Capacity and Utilization collection) in its own region.
- Has no knowledge of the other regions.
- Replicates its data up to the master region.

4.1.1.2. Region Settings

In the **Region** area, set items that apply to your entire Red Hat CloudForms infrastructure such as users, LDAP Groups, capacity and utilization collection, company tags and tag categories, and licensing. Regions are also used for database replication.

4.1.1.3. Capacity and Utilization Collections

4.1.1.3.1. Capacity and Utilization Collection Settings

Use **C & U Collection Settings** to select specifically which clusters and datastores you want to collect usage data for. By selecting a cluster, you are choosing to collect data for all hosts and virtual machines that are part of that cluster. You must also have a server with the Capacity & Utilization **Coordinator**, **Data Collector**, and **Data Processor** roles enabled as well. See Section **Server Control Settings**.

After a provider has been discovered and its relationships refreshed, view the clusters, hosts, and datastores from the settings menu. Navigate to **Configuration**, then click on the **Settings** → **Region** → **C & U Collection** tab.

4.1.1.3.2. Enabling a Cluster, Host, or Datastore for Capacity and Utilization Collection

To enable a cluster, host, or datastore for Capacity and Utilization Collection:

1. From the settings menu, select **Configuration**, then click on the **Settings** accordion.
2. Select **Region**, then click on the **C & U Collection** tab.
3. In the **Clusters** area, check all clusters and hosts that you want to collect data for.
4. In the **Datastores** area, check all datastores that you want to collect data for.
5. Click **Save**.



NOTE

1. As new clusters, hosts, and datastores are discovered, you will need to come back to this configuration to enable collection of capacity and utilization data unless you have used the **Collect for All** check boxes.
2. **Collect for All Clusters** must be checked to be able to collect capacity and utilization data from cloud providers such as Red Hat OpenStack Platform or Amazon EC2.

4.1.1.4. Tags

4.1.1.4.1. Company Tag Categories and Tags

Red Hat CloudForms allows you to create your own set of tags and tag categories. Use tags to create a customized, searchable index for your resources. Depending on your database type, your tags may be case sensitive. After creating these values, you can apply them to your resources. There are two kinds of tags.

- **Company tags** which you will see under **My Company Tags** for a resource. Create company tags from the settings menu. Navigate to **Configuration**, then click the **Settings** accordion, then **Region** → **My Company Tags**. A selection of company tags is provided to you by default as samples. These can be deleted if you do not need them, but are not recreated by Red Hat CloudForms.
- **System tags** are assigned automatically by Red Hat CloudForms.

4.1.1.4.2. Creating a Tag Category

To create a tag category:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Categories** tab.

3. Click **Add** under the **Actions** column to create a new category.

4. In the **Category Information** area:

Category Information

Name	<input type="text"/>
Description	<input type="text"/>
Long Description	<input type="text"/>
Show in Console	<input checked="" type="checkbox"/> ON <input type="checkbox"/>
Single Value	<input checked="" type="checkbox"/> ON <input type="checkbox"/>
Capture C & U Data by Tag	<input type="checkbox"/> OFF

* 'Name' and 'Single Value' fields cannot be edited after adding a category.

- Use **Name** to create a short name that refers to category in the VMDB.




NOTE

The **Name** and **Single Value** fields cannot be changed after the category has been added.

- Use **Description** to type a brief explanation of how the category should be used. This shows when you try to add a value to the category.
- Use **Long Description** to type a detailed explanation of the category.
- Set **Show in Console** to **ON** when the category is ready for use in the console. For example, you want to populate values for the category before exposing it to users.
- Set **Single Value** to **ON** for categories that can only have a single value assigned to a resource. For example, a virtual machine can only be assigned to one location, but could belong to more than one department. This cannot be changed after the category is created.
- Set **Capture C & U Data by Tag** to **ON** for the ability to group capacity and utilization data by this tag category. To use this, be sure to assign this tag to all the resources that you want to group by.

5. Click **Add**.

Repeat these steps for each category you need. After you have created the category, you can add values to it.



IMPORTANT

If no values are created for a category, you are unable to assign a value from that category nor be able to filter by that category.

4.1.1.4.3. Deleting a Tag Category

To delete a tag category:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Categories** tab.
3. Click **Delete** under the **Actions** column for the category you want to delete.
4. Click **OK** to confirm.



NOTE

When you delete a tag category, the category values are removed, and any tags from the category are unassigned from all resources.

4.1.1.4.4. Creating a Company Tag

To create a company tag:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Tags** tab.
3. In the **Choose a Category** area, select a category from the **Category** list.



NOTE

- Some categories only allow one value to be assigned to a resource.
- For some databases such as **PostgreSQL**, tags are case sensitive. For example, filtering by *Linux* in title case give you different results from filtering by *linux* in lower case.

4. Click **Add** under the **Actions** column, and type a **Name** and **Description** for your new value.
5. Click **Add** once again to add the new entry to the table.

4.1.1.4.5. Deleting a Company Tag

To delete a company tag:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Tags** tab.
3. Click **Delete** under the **Actions** column next to the tag to delete it.

**NOTE**

When you delete a tag, the tag is also deleted from any resource to which it was assigned.

4. Click **OK** to confirm.

4.1.1.4.6. Importing Tags for Virtual Machines

You can import a **CSV** file with tag assignments into the VMDB. For the import to be successful, be aware of the following:

- The file must be in the following format, with one line for each virtual machine. One virtual machine per tag must be on a separate line even if you are assigning multiple tags of the same category.
- You must use the display names of the category and the display name for the tag for the import to work.

```
name,category,entry
evm2,Provisioning Scope,All
evm2,Exclusions,Do not Analyze
evm2,EVM Operations,Analysis Successful
rhel6,Department,Presales
rhel6,Department,Support
```

4.1.1.4.7. Importing Tags for a Virtual Machine from a CSV File

To import tags for a virtual machine from a CSV file:

1. Make sure the **CSV file** is in the required format.
2. From the settings menu, select **Configuration**.
3. Click on the **Settings** accordion, then **Region**, then click on the **Import Tags** tab.
4. Click **Choose file** to go to the location where the file is located.
5. Click **Upload**.

**NOTE**

If there are any problems with the file, such as an incorrect column name, unknown virtual machine, unknown tag, or multiple values for a tag that should have only one, an error message will appear in the console for those records.

6. Click **Apply**.

4.1.1.4.8. Importing Custom Values for Virtual Machines and Hosts

You can import a **CSV** file with asset tag information into the VMDB for a virtual machine or import custom values for hosts. For the import to be successful, the file must be in the following format, with one line for each virtual machine or host.

- There are two columns.
- The first line of the file must have the column names as shown below.
- The column names are case sensitive.
- Each value must be separated by a comma.

Virtual Machine Import Example

```
name,custom_1
Ecommerce,665432
Customer,883452
SQLSrvr,1090430
Firewall,8230500
```

For virtual machines, the value for custom_1 will show in the **VM Summary** page as the **Custom Identifier** page as the **Custom Identifier** in the **Properties** area. All of the custom values will show in the **Custom Fields** area.

Host Import Example

```
hostname,custom_1,custom_2
esx303.galaxy.local,15557814,19948399
esxd1.galaxy.local,10885574,16416993
esxd2.galaxy.local,16199125,16569419
```

For hosts, the value for custom_1 will show in the **Host Summary** page as the **Custom Identifier** in the **Properties** area. All of the custom values will show in the **Custom Fields** area.

4.1.1.4.9. Importing Asset Tags for a Virtual Machine from a CSV File

To import asset tags for a virtual machine from a CSV file

1. Make sure the **CSV file** is in the required format.
2. From the settings menu, select **Configuration**.
3. Click on the **Settings** accordion, then **Region**, then click on the **Import** tab.
4. Select the type of custom variable you want to import, either **Host** or **VM**.
5. Click **Choose file** to go to the location where the custom variable file is located.
6. Click **Upload**.



NOTE

If there are any problems with the file, such as an incorrect column name, unknown virtual machine or host, a message appears.

7. Click **Apply**.

4.1.1.5. Registering Red Hat CloudForms

You can register appliances, edit customer information, and apply CloudForms updates from the **Red Hat Updates** tab, accessible from the settings menu, and navigating to **Configuration** → **Region** in the user interface. You can register your appliance to either Red Hat Content Delivery Network (CDN) or to a Red Hat Satellite server, which assign the necessary update packages to the Red Hat CloudForms server. The subscription management service you register with will provide your systems with updates and allow additional management.

The following tools are used during the update process:

- **Yum** provides package installation, updates, and dependency checking.
- **Red Hat Subscription Manager** manages subscriptions and entitlements.
- **Red Hat Satellite Server** provides local system registration and updates from inside the customer's firewall.



IMPORTANT

The update worker synchronizes the VMDB with the status of available Red Hat CloudForms content every 12 hours.



NOTE

Servers with the **RHN Mirror** role also act as a repository for other appliances to pull Red Hat CloudForms package updates.

4.1.1.5.1. Registering Appliances

Before you can access and apply package updates, you must register and subscribe the Red Hat CloudForms appliance to either Red Hat Content Delivery Network (CDN) or to a Red Hat Satellite server.

You need the following to register your appliance:

- Your Red Hat account login or Red Hat Network Satellite login
- A Red Hat subscription that covers your product

To register your appliance with Red Hat Subscription Management or Red Hat Satellite 6, first configure the region with your registration details. These settings will apply to all appliances in this region.

To configure registration for a region:

1. Log in to the appliance as the **admin** user.
2. From the settings menu, select **Configuration**.
3. Select **Region** in the accordion menu and click the **Red Hat Updates** tab.
4. Click **Edit Registration**.
5. Configure registration details for the Red Hat CloudForms appliance using one of two available options:
 - a. To register with Red Hat Subscription Management:

- i. In **Register to**, select **Red Hat Subscription Management**.
 - ii. Enter the **Red Hat Subscription Management Address**. The default is `subscription.rhn.redhat.com`.
 - iii. Enter the **Repository Name(s)**. The default is `cf-me-5.8-for-rhel-7-rpms` `rhel-server-rhsc1-7-rpms`, which are the Red Hat CloudForms repository and the Red Hat Software Collections repository.
 - iv. To use a HTTP proxy, select **Use HTTP Proxy** and enter your proxy details.
 - v. Enter your Red Hat account information and click **Validate**.
 - vi. After your credentials are validated, click **Save**.
- b. To register with Red Hat Satellite 6:
- i. In **Register to**, select **Red Hat Satellite 6**.
 - ii. Enter the **Red Hat Satellite 6 Address**. The default is `subscription.rhn.redhat.com`.
 - iii. Enter the **Repository Name(s)**. The default is `cf-me-5.8-for-rhel-7-rpms` `rhel-server-rhsc1-7-rpms`, which are the Red Hat CloudForms repository and the Red Hat Software Collections repository.
 - iv. To use a HTTP proxy, select **Use HTTP Proxy** and enter your proxy details.
 - v. Enter your Red Hat Satellite account information and click **Validate**.
 - vi. After your credentials are validated, click **Save**.

Your appliance now appears in the **Appliance Updates** list as **Not registered**.

To register your appliance:

1. Select the appliance from the **Appliance Updates** list.
2. Click **Register** to subscribe the appliance and attach subscriptions.

Registering and attaching subscriptions takes a few minutes. The subscription process is complete when the appliance reports that it is **Subscribed** under **Update Status**, and **Registered** under **Last Message**.

You can now apply updates to your appliance.



NOTE

To update your appliances, see [Updating Red Hat CloudForms](#) in *Migrating to Red Hat CloudForms 4.5*.

4.1.1.5.2. Subscription Management for Virtual Environments

Customers can license Red Hat CloudForms for a limited set of providers. This ability is enabled by providing entitlement certificates that describe the features to be enabled. Red Hat CloudForms can be shipped as a bundled product with other Red Hat products like Red

Hat OpenStack Platform and Red Hat OpenShift, providing advanced management capabilities to these products.

Entitlements provides the following enhancements:

- Ability to enable or disable providers based upon a certificate.
- Active subscription with Red Hat Cloud Data Network for delivery to Red Hat CloudForms.
- Ability to remain in its own Red Hat CloudForms channel.
- Ability to add providers even if no certificate is found.
- In the presence of a certificate, providers are limited as per SKU, the certificate is supporting.
- Ability to support the provider to SKU mapping.
- Providers remain fully functional even after adding or removing SKU associated with certificates.



4.1.2. Profiles

4.1.2.1. Creating an Analysis Profile

You can create an analysis profile by referring to the sample profiles provided in the console. You can copy the sample profile or create a new one.

4.1.2.2. Creating a Host Analysis Profile

To create a host analysis profile:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Click  (**Configuration**), and  (**Add Host Analysis Profile**).
4. In the **Basic Information** area, type in a **Name** and **Description** for the analysis profile.

Basic Information

Name

Description


Type

Host

5. Click **File** to collect information about a file or group of files.


6. From the **File Entry** area, click  (**Click to add a new entry**) to add a file or group of files.

File Entry

	Name	Collect Contents?
	<New Entry>	<New Entry>

- Check **Collect Contents** to not only check for existence, but also gather the contents of the file. If you do this, then you can use the contents to create policies in Red Hat CloudForms Control.

7. Click **Event Log** to specify event log entries to collect.

8. From the **Event Log Entry** area, click  (**Click to add a new entry**) to add a type of event log entry. Type in a **Name**. You can type in a specific message to find in **Filter Message**. In **Level**, set the value for the level of the entry and above. Specify the **Source** for the entry. Finally, set the # number of days that you want to collect event log entries for. If you set this to 0, it will go as far back as there is data available.



Event Log Entry

	Name	Filter Message	Level	Source	# of Days
	hostd		warn		5

9. Click **Add**.

4.1.2.3. Creating a Virtual Machine Analysis Profile

To create a virtual machine analysis profile:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Click  (**Configuration**), and  (**Add VM Analysis Profile**).
4. In the **Basic Information** area, type in a **Name** and **Description** for the analysis profile.

Basic Information

Name

Description

Type

Vm

5. You begin in the **Category** tab. From the **Category Selection** area, check the categories you want to collect information for. This is available for virtual machine profiles only.


Category File Registry Event Log

Category Selection

☐ Services ☐ Software ☐ System


☐ User Accounts ☐ VM Configuration

6. Click the **File** tab to collect information about a file or group of files.

7. From the **File Entry** area, click  (**Add this entry**) to add a file or group of files, then type a name. For virtual machines, specify the file to check for. Check the box under **Collect Contents** if you want to collect the file contents as well. The files can be no larger than 1 MB.

Category **File** Registry Event Log

File Entry


	Name	Collect Contents?
	/etc/*.conf	<input checked="" type="checkbox"/>

8. Click the **Registry** tab to collect information on a registry key.

9. From the **Registry Entry** area, type your **Registry Key** and **Registry Value**. To evaluate whether a registry key exists or does not exist on a virtual machine, without providing a value, type * in the **Registry Value** field. Then, you do not need to know the registry value to collect the keys. This is available for virtual machine profiles only.

Category File **Registry** Event Log

Registry Entry

	Registry Hive	Registry Key	Registry Value
	HKLM		*

10. Click **Event Log** to specify event log entries to collect.

11. From the **Event Log Entry** area, complete the fields to add a type of event log entry. You can type in a specific message to find in **Filter Message**. In **Level**, set the value for the level of the entry and above. Specify the **Source for the entry**.

Finally, set the # (number) of days that you want to collect event log entries for. If you set this to 0, it will go as far back as there is data available.

Category File Registry **Event Log**


Event Log Entry

	Name	Filter Message	Level	Source	# of Days
	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

12. Click **Add**.

4.1.2.4. Editing an Analysis Profile


To edit an analysis profile:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Check the analysis profile you want to edit.
4. Click  (**Edit the selected Analysis Profiles**).
5. Make any changes.
6. Click **Save**.

The changes are added to the analysis profile. The virtual machines or hosts must be re-analyzed to collect the new or modified information.

4.1.2.5. Copying an Analysis Profile


To copy an analysis profile:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Check the analysis profile you want to copy.
4. Click  (**Copy the selected Analysis Profiles**).
5. Type a new **Name** and **Description**.
6. Make required changes.
7. Click **Add**.

4.1.2.6. Setting a Default Analysis Profile

If you want to set an analysis profile to be used for all virtual machines, you can create a default profile.

To create a default analysis profile:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Click on the analysis profile you want to set as the default.
4. Click  (**Edit the selected Analysis Profile**).
5. For a virtual machine profile, enter **default** in lower case in **Name**. For a host profile, enter host default.

Basic Information

Name

Description

Type

Vm

6. Click **Save**.

4.1.3. Zones

You can organize your Red Hat CloudForms Infrastructure into zones to configure failover and isolate traffic. A provider that is discovered by a server in a specific zone gets monitored and managed in that zone. All jobs, such as a SmartState Analysis or VM power operation, dispatched by a server in a specific zone can get processed by any Red Hat CloudForms appliance assigned to that same zone.

Zones can be created based on your own environment. You can make zones based on geographic location, network location, or function. When first started, a new server is put into the default zone.

Suppose you have four Red Hat CloudForms appliances with two in the East zone, appliances A and B, and two in the West zone, appliances C and D. VC East is discovered by one of the Red Hat CloudForms appliances in the Red Hat CloudForms Eastern zone. If Appliance A dispatches a job of analyzing twenty virtual machines, this job can be processed by either Appliance A or B, but not C or D.





NOTE

Only users assigned the super administrator role can create zones. There must always be at least one zone. The **Default Zone** is provided and cannot be deleted.

4.1.3.1. Creating a Zone

To create a zone:

1. From the settings menu, select **Configuration**.

- Click on the **Settings** accordion, then click **Zones**.
- Click  (**Configuration**), and  (**Add a new Zone**) to create a zone.
- In the **Zone Information** area, type in a **Name** and **Description** for the new zone.

Zone Information

Name	<input type="text" value="West"/>
Description	<input type="text" value="Western Zone"/>
SmartProxy Server IP	<input type="text" value="192.168.252.12"/>

- Use **SmartProxy Server IP** to specify the IP address of the server that you want SmartProxies installed in this zone to report to. If this is not set, then the IP address of the server that deployed the SmartProxy is used. This does not apply to embedded SmartProxies.
- Optionally, you can configure **NTP servers** for the entire zone in the NTP Servers area. These settings will be used if the NTP servers have not been set for the appliance in the **Operations** → **Server** page.
- In the **Credentials** → **Windows Domain** area, type in Windows domain credentials to be able to collect running processes from Windows virtual machines that are on a domain.

Credentials - Windows Domain

Username	<input type="text" value="acme\admin"/>
Password	<input type="password" value="••••••••"/>
Verify Password	<input type="password" value="••••••••"/>

- In the **Settings** area, set the number for **Max Active VM Scans**. The default is **Unlimited**.
- Click **Save**.



4.1.3.2. Deleting a Zone

To delete a zone:

- From the settings menu, select **Configuration**.
- Click on the **Settings** accordion, then click **Zones**.
- Click the zone you want to remove.



**NOTE**

You cannot delete a zone if there are servers assigned to it.

4. Click  (**Configuration**), then click  (**Delete this Zone**).
5. Click **OK** to confirm.

4.1.3.3. Editing a Zone

To edit a zone:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone you want to edit.
4. Click  (**Configuration**), then click  (**Edit this Zone**).
5. Make the required changes.
6. Click **Save**.

4.1.3.4. Adding SmartProxy Affinity to a Zone:

Enable SmartProxy Affinity for zones containing servers with the SmartProxy role to run a SmartState Analysis.

To add SmartProxy Affinity to a zone:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone in which you want to enable SmartProxy Affinity.
4. Click the **SmartProxy Affinity** tab and click the appropriate server.
5. Click **Save**.

4.1.4. Servers

Server settings enables you to control how each Red Hat CloudForms server operates including authentication, logging, and email. If you have multiple servers in your environment that are reporting to one central VMDB, then you can edit some of these settings from the console by specifying which server you want to change.

**NOTE**

The server selection options are only available if you have multiple servers sharing one VMDB.

4.1.4.1. Changing Server Settings

To change server settings:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the Red Hat CloudForms server is located.
4. In the **Servers** area, click on the Red Hat CloudForms server.
5. Click **Server**.
6. Make any required changes.
7. Click **Save**.

4.1.4.1.1. Basic Information Settings

Server	Authentication	Workers	Custom Logos	Advanced
Basic Information				
Hostname	localhost.localdomain			
IP Address	10.64.15.247			
Resides on VM	🚫 Not Available			
Company Name	<input type="text" value="My Company"/>			
Appliance Name	<input type="text" value="EVM"/>			
Zone*	<input type="text" value="default"/>			
Appliance Time Zone	<input type="text" value="(GMT+00:00) UTC"/>			
Default Locale	<input type="text" value="Client Browser Setting"/>			

* Changing the Zone will reset all of this Server's priorities to secondary.

- Use **Company Name** (maximum 20 characters) to customize the interface with your company's name. You will see the company name when you are viewing or modifying the tags of an infrastructure object or virtual machine.
- Specify the **Appliance Name** (maximum 20 characters) you want displayed as the appliance that you are logged into. You will see this in the upper right corner of the interface with the name of the consoles logged on user.
- Use **Zone** to isolate traffic and provide load balancing capabilities. Specify the zone that you want this Red Hat CloudForms appliance to be a member of. At startup, the zone is set to default.
- Use **Appliance Time Zone** to set the time zone for this server.

**NOTE**

This is the time zone used when created scheduled analyses. This is not the same as the **Time Zone** parameter, which is found by navigating to the settings menu, then **My Settings**, then exploring the **Display Settings** area, and is the time zone displayed in the console.

- Use **Default Locale** to specify the default language for this server.

4.1.4.1.2. Server Control Settings

A server role defines what a server can do. Red Hat recommends that Database Operations, Event Monitor, Reporting, Scheduler, SmartState Analysis, User Interface, Provider Inventory, Provider Operations, and Web Services be enabled on at least one server in each zone. These roles are enabled by default on all servers.

- Use **Default Repository SmartProxy** to set the SmartProxy from which you refresh your virtual machine repositories. This host must have access to your repositories to analyze its virtual machines.

**NOTE**

- Only super administrators can change server roles.
- If you are using more than one Red Hat CloudForms appliance, be sure to set this on all of the appliances.

4.1.4.1.3. Server Roles**NOTE**

- Server roles that are in an active/active high availability configuration (*load balancing and failover protection*) are active in more than one location; whereas, roles that are in an active/passive (*primary/secondary in the case of CloudForms*) high availability configuration (*failover protection*), if more than one CloudForms server in a specific zone or region has this role, only one will be active (*primary*) at a time and a failover has to occur to the passive (*secondary*) appliance with that role.
- For information on region and zone diagnostics and server role priorities, see [Section 4.3, “Diagnostics”](#).

Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
Automation Engine	Use this role if you want to use this CloudForms server to process automation tasks.	N/A	Active/Active

Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
Capacity and Utilization Coordinator	The Capacity & Utilization Coordinator role checks to see if it is time to collect data, somewhat like a scheduler. If it is time, a job is queued for the Capacity and Utilization Data Collector. The coordinator role is required to complete Capacity and Utilization data collection. If more than one CloudForms server in a specific zone has this role, only one will be active at a time.	Zone	Primary/Secondary
Capacity & Utilization Data Collector	The Capacity & Utilization Data Collector performs the actual collection of capacity and utilization data. This role has a dedicated worker, and there can be more than one CloudForms server with this role in a zone.	Zone	Active/Active
Capacity & Utilization Data Processor	The Capacity & Utilization Data Processor processes all of the data collected, allowing CloudForms to create charts. This role has a dedicated worker, and there can be more than one CloudForms server with this role in a zone.	Zone	Active/Active
Database Operations	Use Database Operations to enable this CloudForms server to run database backups or garbage collection.	Zone	Active/Active
Embedded Ansible	This role is disabled by default. The Embedded Ansible role supports Ansible Automation Inside functionality. Enable this role to configure playbook repositories and run playbooks natively to back service catalog items.	Region	Primary/Secondary

Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
Event Monitor	This role is enabled by default and provides the information shown in timelines. The Event Monitor is responsible for the work between the CloudForms server and your providers. It starts 2 workers for each provider. One worker, the monitor, is responsible for maintaining a connection to a provider, catching events, and putting them on the CloudForms message queue for processing. The second worker, the handler, is a message queue worker responsible for delivering only those messages for a provider. You should have at least one of these in each zone.	Zone	Primary/Secondary
Git Repository	The Git Repositories Owner server role supports importing domains into automate from a git repository. This feature is available from the Automate > Import/Export screen in the CloudForms user interface.	Region	Primary/Secondary
Notifier	Use this role if you will be using CloudForms Control or Automate to forward SNMP traps to a monitoring system or send e-mails. See Section 4.1.4.1.6, “Configuring SNMP” for details on creating SNMP alerts. If more than one CloudForms server in a specific region has this role, only one will be active at a time.	Region	Primary/Secondary


Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
Provider Inventory	This role is enabled by default. This role is responsible for refreshing provider information including EMS, hosts, virtual machines, and clusters, and is also responsible for capturing datastore file lists. If more than one CloudForms server in a specific zone has this role, only one will be active at a time.	Zone	Primary/Secondary
Provider Operations	This role is enabled by default. This role sends stop, start, suspend, shutdown guest, clone, reconfigure, and unregister to the provider, directly from the console or through a policy action if you have CloudForms Control. More than one CloudForms server can have this role in a zone.	Zone	Active/Active
RHN Mirror	An appliance with RHN Mirror enabled acts as a server containing a repository with the latest CloudForms packages. This also configures other appliances within the same region to point to the chosen RHN Mirror server for updates. This provides a low bandwidth method to update environments with multiple appliances.	N/A	Active/Active
Reporting	This role is enabled by default. The Reporting role specifies which CloudForms servers can generate reports. If you do not have a CloudForms server set to this role in a zone, then no reports can be generated in that zone. You should have at least one of these in each zone.	Zone	Active/Active

Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
Scheduler	This role is enabled by default. The Scheduler sends messages to start all scheduled activities such as report generation and SmartState analysis. This role also controls all system schedules such as capacity and utilization data gathering. One server in each region must be assigned this role or scheduled CloudForms events will not occur. If more than one CloudForms server in a specific region has this role, only one will be active at a time.	Region	Primary/Secondary
SmartProxy	Enabling the SmartProxy role turns on the embedded SmartProxy on the CloudForms server. The embedded SmartProxy can analyze virtual machines that are registered to a host and templates that are associated with a provider. To provide visibility to repositories, install the SmartProxy on a host from the CloudForms console. This SmartProxy can also analyze virtual machines on the host on which it is installed. Enabling the SmartProxy role on an appliance requires selecting the SmartProxy Affinity for a zone to run a SmartState Analysis. By default, no selections are enabled under SmartProxy Affinity.	Zone	Active/Active
SmartState Analysis	This role is enabled by default. The SmartState Analysis role controls which CloudForms servers can control SmartState Analyses and process the data from the analysis. You should have at least one of these in each zone.	Zone	Active/Active

Server Role	Description	Zone or Region Aware	Primary/Secondary or Active/Active
User Interface	This role is enabled by default. The Web Services role must also be enabled with this role to log into the user interface, as the User Interface role queries the API to receive tokens for login. Uncheck User Interface if you do not want users to be able to access this CloudForms server using the CloudForms console. For example, you may want to turn this off if the CloudForms server is strictly being used for capacity and utilization or reporting generation. More than one CloudForms server can have this role in a zone.	Zone	Active/Active
Web Services	This role is enabled by default. The Web Services role provides API access and must be enabled if the User Interface role is enabled to log into the user interface. You can also enable the Web Services role to provide API-only access to the server. Uncheck Web Services to stop this CloudForms server from acting as a web service provider. More than one CloudForms server can have this role in a zone.	N/A	Active/Active
Websocket	This role enables starting or stopping websocket workers required for proxying remote consoles.	N/A	Active/Active

4.1.4.1.4. VMware Console Settings

If you are using the Red Hat CloudForms control feature set, then you have the ability to connect to a Web console for virtual machines that are registered to a host. To use this feature, you must have VNC installed, [VMware's WebMKS SDK enabled in CloudForms](#), or the VMRC native desktop application installed for your environment.

**NOTE**

See vendor documentation to ensure you are installing appropriate applications for your virtual infrastructure. Once you have installed the required software, you must specify its version in the CloudForms configuration settings.

VMware Console Support

Use	<input type="text" value="VNC"/>
VNC Proxy Address	<input type="text"/>
VNC Proxy Port	<input type="text"/>

- If you select **VNC**, type in the port number used. This port must be open on the target virtual machine and the VNC software must be installed there. On the computer that you are running the console from, you must install the appropriate version of Java Runtime if it is not already installed.
- If you select **VMware WebMKS**, select the appropriate version.
- If using **VMware VMRC** desktop application, be sure that you have fulfilled the requirements for your vCenter version. The correct version of the VMRC desktop application from VMware must be installed on the client computer. To do this, log into the vCenter Web Service and attempt to open a virtual machine console. The vSphere Web Client must be installed on vCenter version 5, and the provider must be registered to it. For vCenter version 4, the VMware vCenter Management Webservices must be running.

4.1.4.1.5. NTP Servers Settings

In the **NTP Servers** area, you can specify the NTP servers to use as source for clock synchronization here. The NTP settings specified here will override Zone NTP settings. Enter one NTP server hostname or IP address in each text box.

4.1.4.1.6. Configuring SNMP

You can use Simple Network Management Protocol (SNMP) traps to send alerts for various aspects of a Red Hat CloudForms environment.

Requirements

- Configure your SNMP management station to accept traps from Red Hat CloudForms appliances. Consult your management station's documentation.
- Each appliance that could process SNMP traps must have the **snmpd** and **snmptrapd** daemons running.
- The region where the appliances are located must have the **Notifier** role enabled and the failover role priority set.

To enable the **snmpd** and **snmptrapd** daemons

1. Access each SNMP processing appliance using SSH.
2. Set the SNMP daemons to run on start up:

```
# chkconfig --level 2345 snmpd on
# chkconfig --level 2345 snmptrapd on
```



3. The daemons run automatically when the appliance is restarted, but must be started manually now.

```
# service snmpd start
# service snmptrapd start
```

To enable the notifier role:


1. Access each SNMP processing appliance using their web interfaces.
2. From the settings menu, select **Configuration → Settings**.
3. Select the zone where the EVM server is located, and select the EVM server.
4. In the **Server Control** area, set the **Notifier** server role option to **ON**.
5. Click **Save**.

To set the failover priority role:

1. From the settings menu, select **Configuration → Diagnostics**.
2. Select the zone where the EVM server is located.
3. Click **Roles by Servers** or **Servers by Roles** to view your servers.
4. In the **Status of Roles for Servers** in **Zone Default** Zone area, click the role that you want to set the priority for.
5. Click  (**Configuration**), and  (**Promote Server**) to make this the primary server for this role.

4.1.4.1.7. Outgoing SMTP Email Settings

To use the email action in Red Hat CloudForms, set an email address to send emails from.



NOTE

To be able to send any emails from the server, you must have the Notifier server role enabled. You can test the settings without the role enabled.

Outgoing SMTP E-mail Server

Host	<input type="text" value="localhost"/>
Port	<input type="text" value="25"/>
Domain	<input type="text" value="mydomain.com"/>
Start TLS Automatically	<input checked="" type="checkbox"/> ON
SSL Verify Mode	<input type="text" value="None"/>
Authentication	<input type="text" value="login"/>
User Name	<input type="text" value="evmadmin"/>
Password	<input type="password"/>
From E-mail Address	<input type="text" value="cfadmin@cfserver.com"/>
Test E-mail Address	<input type="text"/> <input type="button" value="Verify"/>

- Use **Host** to specify the host name of the mail server.
- Use **Port** to specify the port for the mail server.
- Use **Domain** to specify the domain name for the mail server.
- Set **Start TLS Automatically** on **ON** if the mail server requires TLS.
- Select the appropriate **SSL Verify Mode**.
- Use the **Authentication** drop down to specify if you want to use **login**, **plain**, or no authentication.
- Use **User Name** to specify the user name required for login authentication.
- Use **From E-mail Address** to set the address you want to send the email from.
- Use **Test E-mail Address** if you want to test your email settings. Click **Verify** to send a test email.

4.1.4.1.8. Web Services Settings

Web services are used by the server to communicate with the SmartProxy.

Web Services

Mode	<input type="text" value="invoke"/>
Security	<input type="text" value="none"/>

- Set **Mode** to **invoke** to enable 2-way Web services communication between the Red Hat CloudForms appliance and the SmartProxy. Set **Mode** to **disabled** to use Web services from the SmartProxy to the Red Hat CloudForms appliance only. When the Red Hat CloudForms appliance has work for the SmartProxy, the work will be placed in a queue in the VMDB. The work will be completed either when the Red Hat CloudForms appliance is able to contact the SmartProxy or when the next SmartProxy heartbeat occurs, whichever comes first.

- If **Web Services** are enabled, you have the option to use **aws-security**.

4.1.4.1.9. Logging Settings

Logging

Log Level

- Use **Log Level** to set the level of detail you want in the log. You can select from **fatal**, **error**, **warn**, **info**, and **debug**. The default setting is **info**.

4.1.4.1.10. Custom Support URL Settings

Custom Support URL

URL (i.e. www.mypage.com)

Description

- Use **URL** to specify a specific URL that you want to be accessible from the **About Product Assistance** area.
- Use **Description** to set a label for the **URL**.

4.1.4.2. Authentication

Use the **Authentication** tab to specify how you want users authenticated on the console. You can use the VMDB or integrate with LDAP, LDAPS, Amazon, or an external IPA server.

4.1.4.2.1. Changing Authentication Settings

To change authentication settings:

1. From the settings menu, select **Configuration**.
2. Click the **Settings** accordion, then click **Zones**.
3. Click the zone where the server is located.
4. Click the server.
5. Click the **Authentication** tab.
6. Use **Session Timeout** to set the period of inactivity before a user is logged out of the console.
7. Set the authentication method in **Mode**. Choose from **Database** (using the VMDB), **LDAP** (Lightweight Directory Authentication Protocol), **LDAPS** (Secure Lightweight Directory Authentication Protocol), **Amazon**, or **External (httpd)**. The default is

Database.

- If you choose **Database**, see [Section 4.2.5, “Creating a User”](#) for steps to create users.
- If you choose **LDAP** or **LDAPS**, see [Section 4.1.4.2.2, “Configuring LDAP Authentication”](#) for configuration steps.
- If you choose **Amazon**, see [Section 4.1.4.2.9, “Amazon Settings”](#) for configuration steps.
- If you choose **External (httpd)**, see [Section 4.1.4.2.10, “External Authentication \(httpd\)”](#) for configuration steps.

8. Click **Save**.

4.1.4.2.2. Configuring LDAP Authentication

If you choose LDAP or LDAPS as your authentication mode, the required parameters are exposed under **LDAP Settings**. Be sure to validate your settings before saving them.

To configure CloudForms to use LDAP for authentication, complete the following steps:

1. From the settings menu, select **Configuration**. Select the **Authentication** tab.
2. Select a **Session Timeout** to set the period of inactivity before a user is logged out of the console.
3. Select **LDAP** or **LDAPS** from the **Mode** list. This exposes additional required parameters under **LDAP Settings**.
4. Configure your **LDAP Settings** (the following example configures a Red Hat Identity Management LDAP server):
 - Use **LDAP Host Names** to specify the fully qualified domain names of your LDAP servers. CloudForms will search each host name in order until it finds one that authenticates the user. Note, CloudForms supports using a maximum of 3 possible **LDAP Host Names**.
 - Use **LDAP Port** to specify the port for your LDAP server. The default is 389 for LDAP and 636 for LDAPS.
 - From the **User Type** list, select one of the following and configure the values for your LDAP server:
 - **User Principal Name**: Type the user name in the format of `user@domainname`, for example, `dbright@acme.com`. (In this case, the user would log on as `dbright`.)
 - **Email Address**: Logs in with the user’s email address.
 - **Distinguished Name (CN=<user>)**: Uses the common name for the user. Be sure to enter the correct **User Suffix** and **Distinguished Name** option for your directory service implementation: for example, `cn=dan bright,ou=users,dc=acme,dc=com`. (The user logs on as `dan bright`.)
 - **Distinguished Name (UID=<user>)**: Uses the user ID (UID). Be sure to enter the correct **User Suffix** and **Distinguished Name** option for your

directory service implementation: for example, **uid=dan bright,ou=users,dc=acme,dc=com**. (The user logs on as **dan bright**.)

- **SAM Account Name:** User login for Active Directory clients and servers using legacy Windows versions. You must also configure the **Domain Prefix** in the next field when selecting this option.
- Specify the **Domain Prefix** if you are configuring an Active Directory LDAP host, and selected **SAM Account Name** as the **User Type**. This field represents the prefix name in the Active Directory domain. Together with the SAM account name, this constructs the fully qualified user name in the format **<domain_prefix>\<user>**.
- Specify the **User Suffix**, such as **acme.com** for **User Principal Name** or **cn=users,dc=acme,dc=com** for **Distinguished Name**, in **Base DN**.



NOTE

The **ldapsearch(1)** command can be used to get details of your LDAP settings. To get details related to a specific user, run:

```
# ldapsearch -D "cn=directory manager" -H
ldap://www.acme.com:389 -b "dc=acme,dc=com" -s sub "
(objectclass=*)" -w password | grep -i dbright
```

To search for your LDAP server's distinguished name (DN) values, run:

```
# ldapsearch -D "cn=directory manager" -H
ldap://www.acme.com:389 -b "dc=acme,dc=com" -s sub "
(objectclass=*)" -w password
```

Example: LDAP Configuration

Settings Server "EVM [1000000000001]" (current)

Server **Authentication** Workers Custom Logos Advanced

Authentication

Session Timeout h m

Mode **LDAP**

LDAP Settings

LDAP Host Names

LDAP Port

User Type **User Principal Name**

Domain Prefix: <domain>\
 <user>

User Suffix: <user>@

Role Settings

Get User Groups from LDAP ☒

Get Roles from Home Forest ☒

Follow Referrals ☒


Base DN

Bind DN

Bind Password

[Validate](#)

Trusted Forest Settings

	LDAP Hostname	Mode	LDAP Port	Base DN	Bind DN	Bind Password
	<Click on this row to create a new forest>					

[Save](#) [Reset](#)

5. Configure your **Role Settings**: In both LDAP and LDAPS, you can use groups from your directory service to set the role for the authenticated LDAP user. The LDAP user must be assigned one of the account role groups. See [Section 4.1.4.2.4, “Assigning CloudForms Account Roles Using LDAP Groups”](#) for more information.

- For LDAP users not belonging to a group:
 - Select a CloudForms group from the **Default Group for Users** list. This default group can be used for all LDAP users who use LDAP for authentication only. Do not select **Get User Groups from LDAP**, which will hide the **Default Group for Users** option.
- For LDAP users belonging to a group:
 - Check **Get User Groups from LDAP** to retrieve the user’s group membership from LDAP. This is used for mapping a user’s authorization to a CloudForms role. This requires group names on the LDAP server to match CloudForms group names.



IMPORTANT

If you do not check **Get User Groups from LDAP**, the user must be defined in the VMDB using the console where the User ID is the same as the user’s name in your directory service typed in lowercase.

- Check **Get Roles from Home Forest** to use the LDAP roles from the LDAP user's home forest. This will allow you to discover groups on your LDAP server and create CloudForms groups based on your LDAP server's group names. Any user logging in will be assigned to that group. This option is only displayed when **Get User Groups from LDAP** is checked.
 - Check **Follow Referrals** to look up and bind a user that exists in a domain other than the one configured in the LDAP authentication settings.
 - Specify the user name to bind to the LDAP server in **Bind DN**. This user must have read access to all users and groups that will be used for CloudForms authentication and role assignment, for example, a service account user with access to all LDAP users.
 - Enter the password for the Bind DN user in **Bind Password**.
6. Click **Validate** to verify your settings.
 7. Click **Save**.

LDAP authentication is now configured in your CloudForms environment.

To assign account roles using LDAP groups, see [Section 4.1.4.2.4, "Assigning CloudForms Account Roles Using LDAP Groups"](#).


4.1.4.2.3. Trusted Forests

Optionally, if a user has group memberships in another LDAP forest, specify the settings to access the memberships in the trusted forest.

When trusted forests are added to the authentication configuration, they are used only for finding groups that a user is a member of. CloudForms will first collect all of the user's groups from the primary LDAP directory. Then it will collect any additional groups that the user is a member of from all of the configured forests.

The collected LDAP groups are used to match, by name, against the groups defined in CloudForms. The user must be a member of at least one matching LDAP group to be successfully authenticated.

To add settings for a trusted forest:

1. From the settings menu, select **Configuration**.
2. Click the **Settings** accordion, then click **Zones**.
3. Click the **Zone** where the server is located.
4. Click the **Server**.
5. Click the **Authentication** tab.
6. Check **Get User Groups from LDAP**, and enter all items in the **Role Settings** area.
7. In the **Trusted Forest Settings** area, click  (**Click to add a new forest**).

8. Enter the **LDAP Host Name**, select a **Mode**, and enter an **LDAP Port**, **Base DN**, **Bind DN**, and **Bind Password**.
9. Click **Save**.

After adding other trusted LDAP forests, you can then change the order in which CloudForms looks up the forests for authentication. For instructions, see [Section 4.1.4.2.7, “Configuring Lookup Priority for LDAP Groups”](#).

4.1.4.2.4. Assigning CloudForms Account Roles Using LDAP Groups

After configuring LDAP authentication as described in [Section 4.1.4.2.2, “Configuring LDAP Authentication”](#), you can associate CloudForms account roles with your LDAP users. The LDAP server defines the groups and users for CloudForms, while CloudForms defines the account roles, and maps the roles to the privileges the LDAP user has.

There are two ways to associate your LDAP groups with CloudForms account roles:

- Create groups in CloudForms that match your existing LDAP groups by name, and assign the groups account roles; or
- Create groups on your LDAP server based on the default account roles in CloudForms.



The users in your LDAP groups then inherit the CloudForms account roles for the LDAP group(s) they are in.

The authentication process then happens as such:

1. **LDAPUser1** attempts to log into CloudForms, so CloudForms queries the LDAP server to verify it knows **LDAPUser1**.
2. The LDAP server then confirms that it knows **LDAPUser1**, and provides information about the LDAP groups **LDAPUser1** belongs to: **Group1**.
3. CloudForms then looks up **Group1**, and discovers that **Group1** is associated with **Role1**.
4. CloudForms then associates **LDAPUser1** with **Group1** in CloudForms, and then allows the user to perform tasks allowable by that role.

4.1.4.2.5. Using Existing LDAP Groups to Assign Account Roles

This section provides instructions for mapping your existing LDAP groups to account roles in CloudForms. As a result, the users in the LDAP group will then be assigned to the CFME roles associated with that group.

1. From the settings menu, select **Configuration**.
2. Click the **Access Control** accordion, then click **Groups**.
3. Click  (**Configuration**), and  (**Add a new Group**) to create a group.
4. There are two ways to specify the group to use:

- In the **Description** field, enter the (common name **cn**) for your existing LDAP group assigned to users requiring access to CloudForms.
 - Select **Look Up LDAP Groups** to find a list of groups assigned to a specific user in LDAP, then use the **LDAP Group for User** list to choose a group.
 - a. In **User to Look Up**, enter the common name (**cn**) for a user in your LDAP group.
 - b. Enter the **Username**.
 - c. In **Password**, enter the user's password. Click **Retrieve**.
5. Select a **Role** to map to the group.
 6. Select a **Project/Tenant** to map to the group.

Adding a new Group

Group Information

Description	<input type="text" value="idm-users"/> <input checked="" type="checkbox"/> (Look Up LDAP Groups)
Role	<input type="text" value="EvmRole-administrator"/>
Project/Tenant	<input type="text" value="My Company"/>

LDAP Group Look Up

User to Look Up	<input type="text" value="cn=idm-users,cn=groups,cn=accounts,dc=lab,dc=local"/>
Username	<input type="text" value="uid=dbright,cn=users,cn=accounts,dc=lab,dc=local"/>
Password	<input type="password" value="....."/>

[Retrieve](#)

7. Select any filters to apply to what this group can view in the **Assign Filters** area:
 - a. In the **My Company Tags** tab, select tags to limit the user to items containing those tags. The items that have changed show in a blue italicized font.
 - b. In the **Host & Clusters** tab, select the host and clusters to limit the user to. The items that have changed show in a blue italicized font.

Assign Filters

This user is limited to the selected items and their children.



- c. In the **VMs & Templates** tab, select the folders created in your virtual infrastructure to limit the user to. The items that have changed show in a blue italicized font.

8. Click **Add**.

To configure the LDAP group lookup priority, see [Section 4.1.4.2.7, “Configuring Lookup Priority for LDAP Groups”](#).

4.1.4.2.6. Using Groups Named by CloudForms to Assign Account Roles

You can also configure access control for LDAP users by creating groups on your LDAP server based on CloudForms account roles.

Your LDAP group names must match the account role names in CloudForms. The LDAP users in that group are then automatically assigned to that specific account role.

1. In your LDAP directory service, define a distribution group for one or more of the account roles with the names shown in the table below. This group must be in the LDAP directory source you specified for the server. See [Section 4.1.4.2.2, “Configuring LDAP Authentication”](#).

Directory Service Distribution Group Name	Account Role
EvmGroup-administrator	Administrator
EvmGroup-approver	Approver
EvmGroup-auditor	Auditor
EvmGroup-consumption_administrator	Consumption Administrator
EvmGroup-container_administrator	Container Administrator
EvmGroup-container_operator	Container Operator
EvmGroup-desktop	Desktop
EvmGroup-operator	Operator
EvmGroup-security	Security
EvmGroup-super_administrator	Super Administrator
EvmGroup-support	Support
EvmRole-tenant_administrator	Tenant Administrator
EvmRole-tenant_quota_administrator	Tenant Quota Administrator
EvmGroup-user	User
EvmGroup-user_limited_self_service	User Limited Self Service

Directory Service Distribution Group Name	Account Role
EvmGroup-user_self_service	User Self Service
EvmGroup-vm_user	Vm User

2. Make each user of your directory service that you want to have access to CloudForms a member of one of these groups.
3. From the settings menu, select **Configuration**.
4. Click the **Settings** accordion, then select your server under **Zones**.
5. Click the **Authentication** tab and enable **Get User Groups from LDAP** after typing in all of the required LDAP authentication settings. See [Section 4.1.4.2.2, “Configuring LDAP Authentication”](#).



4.1.4.2.7. Configuring Lookup Priority for LDAP Groups

CloudForms can have multiple LDAP groups configured, which the appliance will attempt to authenticate with one by one until it succeeds. The lookup priority of these groups can be rearranged; to configure the order in which CloudForms looks up LDAP groups:



NOTE

On initial login, a user’s *current group* assignment is the highest priority group. User group membership, on subsequent logins, is set as the last assigned group from the prior session.

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Groups**.
3. Click  (**Configuration**), and  (**Edit Sequence of User Groups for LDAP Look Up**) to prioritize which group a user will default to if LDAP returns multiple matching groups.
4. Select one or more consecutive groups and use the arrow buttons to move the user group higher or lower in priority.
5. Click **Save**.

4.1.4.2.8. Testing LDAP Configuration

To test that your LDAP or LDAPS group configuration is working correctly with CloudForms:

1. Log out of the CloudForms user interface.
2. Log back in as an LDAP user that is assigned to one or more of the matching groups.

3. Change groups by clicking on the user dropdown menu on the top right of the user interface. The dropdown list will show the groups the user is authorized for.

You can also check the logs in `/var/www/miq/vmdb/log/audit.log` or `/var/www/miq/vmdb/log/evm.log` to verify your LDAP configuration is working correctly.

First, run the following command in a terminal to view the log messages in real time:

```
tail -f /var/www/miq/vmdb/log/audit.log
```

Then log into the CloudForms user interface as an LDAP user, while checking `/var/www/miq/vmdb/log/audit.log` for updated status, success, or failure messages. Alternatively, you can test your LDAP configuration by viewing the logs in `/var/www/miq/vmdb/log/evm.log` with `grep`, which are more verbose.

Troubleshooting LDAP configuration

To test a problematic CloudForms LDAP configuration, run the following to see if the user been pulled from LDAP with the right group. For example:

```
# ldapsearch -x -H ldap://ldap-example:389 -LLL \ -b
"ou=people,dc=example,dc=com" -s sub \ -D
"ui=:userid,ou=People,dc=example,dc=com" -w :password \ "
(objectclass=organizationalPerson)
```

To test if the user belongs to right group, include one of the following lines in the `ldapsearch` command:

```
(&(objectClass=user)(sAMAccountName=yourUserName)
(memberof=CN=YourGroup,OU=Users,DC=YourDomain,DC=com))
```

or

```
-b "ou=groups, dc=example,dc=com"
```

4.1.4.2.9. Amazon Settings

If you choose Amazon as your authentication mode, required parameters are exposed under **Amazon Primary AWS Account Settings** for **IAM**. Be sure to validate your setting before saving them.

- Type in an **Access Key** provided by your Amazon account.
- Type in a **Secret Key** provided by your Amazon account.

Users logging into CloudForms with Amazon authentication enter their own IAM Access Key as the username and IAM Secret Key as the password. Amazon users must be added as a CloudForms user or belong to an IAM user group added to the list of CloudForms groups.

4.1.4.2.10. External Authentication (httpd)

When external authentication is enabled, users can log in to the CloudForms appliance using their IPA server credentials. The appliance creates user accounts automatically and imports relevant information from the IPA Server.

The appliance contains IPA client software for connecting to IPA servers, but it is not configured by default. External authentication is enabled by first configuring it in the web interface, then in the console. Disabling external authentication and returning to internal database authentication also requires steps in both the web interface and the console.

Requirements

- For an appliance to leverage an IPA Server on the network, both the appliance and the IPA server must have their clocks synchronized or Kerberos and LDAP authentication fail.
- The IPA Server must be known by DNS and accessible by name. If DNS is not configured accordingly, the hosts files need to be updated to reflect both IPA server and the appliance on both virtual machines.
- For users to log in to the appliance using IPA server credentials, they must be members of at least one group on the IPA server which is also defined in the appliance. Navigate to the settings menu, then **Configuration** → **Access Control** → **Groups** to administer groups.

Configuring the Appliance for External Authentication

To configure the appliance for external authentication, first set up authentication using the web interface, then using the console.

Using the web interface:

1. Log in to the web interface as an administrative user.
2. Navigate to the settings menu, then **Configuration** → **Access Control** → **Zone** → **Server** → **NTP Servers** or use the hosting provider of the virtual machine to synchronize the appliance's time with an NTP server.
3. Navigate to the settings menu, then **Configuration** → **Authentication**.
4. Select a **Session Timeout** if required.
5. Select **External (httpd)** in the **Mode list**.
6. Select **Enable Single Sign-On** to allow single sign-on using Kerberos tickets from client machines that authenticate to the same IPA server as the appliance.
7. In the **Role Settings** area, select **Get User Groups** from **External Authentication (https)**.
8. Click **Save**.

Using the console:

1. Log in to the appliance console using the user name **admin**.
2. The summary screen displays:

```
| External Auth:  not configured
```

3. Press Enter.

4. Enter **10** to select Configure External Authentication (httpd).
5. Enter the fully qualified hostname of the IPA Server, for example *ipaserver.test.company.com*.
6. Enter the IPA server domain, for example *test.company.com*.
7. Enter the IPA server realm, for example *TEST.COMPANY.COM*.
8. Enter the IPA server principal, for example *admin*.
9. Enter the password of the IPA server principal.
10. Enter **y** to proceed.



NOTE

If any of the following conditions are true, configuration fails:

- The IPA server is not reachable by its FQDN
- The IPA server cannot reach the appliance by its FQDN
- The time is not synchronized between the appliance and the IPA server
- The IPA server admin password is entered incorrectly

Reverting to Internal Database Authentication

To revert to internal database authentication, first configure authentication using the web interface, then using the console.

Using the web interface:

1. Log in to the web interface as an administrative user.
2. Navigate to the settings menu, then **Configuration → Authentication**.
3. Select **Database** in the Mode list.
4. Click **Save**.

Using the console:

1. Log in to the appliance console using the user name **admin**.
2. The summary screen displays:

```
External Auth:  IPA.server.FQDN
```

3. Press **Enter**.
4. Enter **10** to select Configure External Authentication (httpd). The currently configured IPA server hostname and domain are displayed.
5. Enter **y** to un-configure the IPA client.

Optional Configuration Using the Appliance Console CLI

In addition to using the appliance console, external authentication can optionally be configured and un-configured using the appliance console command line interface.

Appliance console CLI command and relevant options include:

```
/bin/appliance_console_cli --host <appliance_fqdn>
                             --ipaserver <ipa_server_fqdn>
                             --iparealm <realm_of_ipa_server>
                             --ipaprincipal <ipa_server_principal>
                             --ipapassword <ipa_server_password>
                             --uninstall-ipa
```

***--host*:**

updates the hostname of the appliance. If you performed this step using the console and made the necessary updates made to **/etc/hosts** if DNS is not properly configured, you can omit the **--host** option.

***--iparealm*:**

if omitted, the **iparealm** is based on the domain name of the **ipaserver**.

***--ipaprincipal*:**

if omitted, defaults to admin.

Example 4.1. Configuring External Authentication

```
$ ssh root@appliance.test.company.com
[appliance]# /bin/appliance_console_cli --host
appliance.test.company.com \
                                     --ipaserver
ipaserver.test.company.com \
                                     --iparealm TEST.COMPANY.COM \
                                     --ipaprincipal admin \
                                     --ipapassword smartvm1
```

Example 4.2. Reverting to Internal Database Authentication

```
$ ssh root@appliance.test.company.com
[appliance]# /bin/appliance_console_cli --uninstall-ipa
```

4.1.4.2.11. Configuring External Authentication Using SAML

This procedure outlines how to manually configure an appliance to use SAML external authentication. While other SAML identity providers can be used with CloudForms, this procedure covers using Red Hat Single Sign-On (SSO) 7.0, which is implemented using the Apache HTTP server's **mod_auth_mellon** module.

To enable external authentication using SAML, complete the following steps to configure your HTTP server, then your CloudForms appliance.



NOTE

The current SAML implementation only secures the CloudForms appliance's web administrative user interface with SAML. The REST API and self service user interface do not currently support SAML.

Requirements

The following is required in order to enable SAML authentication to the appliance:

- A CloudForms 4.5 appliance
- A SAML identity provider (e.g. Red Hat Single Sign-On (SSO) 7.0 or later)

Configuring the HTTP Server for SAML

The Apache HTTP server first must be configured to work with SAML authentication. All SAML-related certificates and keys are accessed from **/etc/httpd/saml2/**.

1. Log into the CloudForms appliance as root using SSH, and create the **/etc/httpd/saml2/** directory:

```
# mkdir -p /etc/httpd/saml2
```

2. Copy the **httpd** remote user and SAML template configuration files to the appliance:

```
# TEMPLATE_DIR="/opt/rh/cfme-appliance/TEMPLATE"
# cp ${TEMPLATE_DIR}/etc/httpd/conf.d/manageiq-remote-user.conf
  /etc/httpd/conf.d/
# cp ${TEMPLATE_DIR}/etc/httpd/conf.d/manageiq-external-auth-
  saml.conf /etc/httpd/conf.d/
```



NOTE

The following are notable SAML configuration defaults in the **manageiq-external-auth-saml.conf** file:

- Identity Provider Files (i.e. Red Hat SSO)
 - Metadata File: **/etc/httpd/saml2/idp-metadata.xml**
- Service Provider Files (i.e. **mod_auth_mellon**)
 - Private Key File: **/etc/httpd/saml2/miqsp-key.key**
 - Certificate File: **/etc/httpd/saml2/miqsp-cert.cert**
 - Metadata File: **/etc/httpd/saml2/miqsp-metadata.xml**

Other **mod_auth_mellon** parameters, such as endpoints and protected URLs, must not be modified as the appliance expects them to be defined as such.

3. Generate the service provider files on the appliance using the Apache HTTP server's **mod_auth_mellon** command **mellon_create_metadata.sh**:

```
# cd /etc/httpd/saml2
# /usr/libexec/mod_auth_mellon/mellon_create_metadata.sh
https://<miq-appliance> https://<miq-appliance>/saml2
```

The **mellon_create_metadata.sh** script creates file names based on the appliance URL.

4. Rename the files created by the **mellon_create_metadata.sh** script to match the expected file names from the **manageiq-external-auth-saml.conf** file:

```
# mv https_<miq-appliance>.key miqsp-key.key
# mv https_<miq-appliance>.cert miqsp-cert.cert
# mv https_<miq-appliance>.xml miqsp-metadata.xml
```

5. Now that the service provider's **metadata.xml** file has been generated, the service provider definition can be defined in the SAML identity provider. For Red Hat SSO, a realm can be created for one or more appliances with individual clients defined one per appliance, where the client ID is specified as the URL of the appliance.

To add a client in the Red Hat SSO CloudForms realm:

- a. Select and import the **miqsp-metadata.xml** file created for **mod_auth_mellon**.
- b. Set the client ID as **https://<miq-appliance>**.
- c. Set the client protocol as **saml**.

6. Update the client definition for the appliance in Red Hat SSO with the following:

Setting	Value
Name ID Format	username
Valid Redirect URIs	<a href="https://<miq-appliance>/saml2/postResponse">https://<miq-appliance>/saml2/postResponse
Assertion Consumer Service POST Binding URL	<a href="https://<miq-appliance>/saml2/postResponse">https://<miq-appliance>/saml2/postResponse
Logout Service Redirect Binding URL	<a href="https://<miq-appliance>/saml2/logout">https://<miq-appliance>/saml2/logout

7. Obtain the identity provider's **idp-metadata.xml** metadata file as follows:

```
# cd /etc/httpd/saml2
# curl -s -o idp-metadata.xml \
  http://<redhatSSO-server>:8080/auth/realms/<miq-
  realm>/protocol/saml/descriptor
```

8. In CloudForms 4.5, the following change is necessary to the **idp-metadata.xml** file for SAML logout to work between **mod_auth_mellon** and Red Hat SSO:

```
# vi idp-metadata.xml

...
<SingleLogoutService
<   Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
---
>   Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
    Location=
...

```

9. Restart the HTTP server on the appliance:

```
# systemctl restart httpd
```

Configuring the Appliance Administrative User Interface

After configuring the HTTP server for SAML, update the CloudForms appliance so that the administrative user interface works with SAML authentication.

1. Login to the appliance as **admin**, and navigate to the settings menu, then **Configuration** → **Authentication**.
2. Set the mode to **External (httpd)**.
3. Check **Enable SAML**. This enables the SAML login button on the appliance login screen, then redirects to the SAML protected page for authentication, and supports the SAML logout process.
4. Check **Enable Single Signon**. With this option enabled, initial access to the appliance's administrative user interface redirects to the SAML identity provider authentication screen. Logging out from the appliance returns the user to the appliance login screen, allowing them to log in as **admin** unless **Disable Local Login** is also checked.
5. Optional: Check **Disable Local Login** to disable the **admin** login to appliance and only allow SAML based authentication. Note that if there are issues with the identity provider or you require **admin** access to the appliance, you cannot log in through the appliance login screen until you re-enable local login as described in [Section 4.1.4.2.11, "Configuring External Authentication Using SAML"](#).
6. Check **Get User Groups from External Authentication (httpd)**
7. Click **Save**.



IMPORTANT

Ensure the user's groups are created on the appliance and appropriate roles are assigned to those groups. See *SAML Assertions* in [Section 4.1.4.2.11](#), “[Configuring External Authentication Using SAML](#)” for more information on the parameters used by the CloudForms appliance.

For example, to configure user groups from your SAML identity provider to work with CloudForms:

1. In your SAML identity provider, specify your existing user groups in similar format to the following:
REMOTE_USER_GROUPS=Administrators;CloudAdministrators;Users
2. On your CloudForms appliance, create the equivalent groups. See *Creating a User Group* in [Section 4.2.8](#), “[Creating a Group](#)”.
3. On your CloudForms appliance, assign EVM roles to the groups. See *Creating a Role* in [Section 4.2.10](#), “[Creating a Role](#)”.

Complete the above steps on each appliance in the settings menu, then navigate to **Configuration → Access Control**.

You can now log into your CloudForms appliance using your SAML credentials.

SAML Assertions

To authenticate to the CloudForms appliance using SAML, the following remote user parameters are looked at by the appliance upon a successful login and redirect from the identity provider. These parameters are used by the appliance to obtain group authentication information.

HTTP Environment	SAML Assertion
REMOTE_USER	username
REMOTE_USER_EMAIL	email
REMOTE_USER_FIRSTNAME	firstname
REMOTE_USER_LASTNAME	lastname
REMOTE_USER_FULLNAME	fullname
REMOTE_USER_GROUPS	groups

For Red Hat SSO, the above SAML assertions can be defined for the appliance client in Red Hat SSO as mappers.

Name	Category	Type	Property
username	AttributeStatement Mapper	User Property	username
email	AttributeStatement Mapper	User Property	email
firstname	AttributeStatement Mapper	User Property	firstName
lastname	AttributeStatement Mapper	User Property	lastName
fullname	AttributeStatement Mapper	User Attribute	fullName
groups	Group Mapper	Group List	groups



IMPORTANT

The **fullName** attribute was not available in the default database as of this writing and was added as a user attribute.

Re-enabling Local Login (Optional)

If you disabled local login in the administrative user interface but need the ability to log in as **admin**, local login can be re-enabled using one of the following methods:

Re-enabling Local Login from the Appliance Administrative User Interface

This method requires the identity provider to be available, and the ability to login as a user with enough administrative privileges to update CloudForms authentication settings.

1. Log in to the appliance user interface as the administrative user.
2. From the settings menu, select **Configuration** → **Authentication**.
3. Uncheck **Disable Local Login**.
4. Click **Save**.

Re-enabling Local Login from the Appliance Console:

1. Use SSH to log into the appliance as **root**.
2. Run the **appliance_console** command.
3. Select **Update External Authentication Options**.
4. Select **Enable Local Login**.
5. Apply the updates.

Alternatively, log into the appliance as root using SSH, and run the following command:

```
# appliance_console_cli --extauth-opts local_login_disabled=false
```

4.1.4.3. Workers

Use the Workers page to specify the number of workers and amount of memory allowed to be used for each type.



NOTE

Only make these changes when directed to by Red Hat Support.

4.1.4.3.1. Changing Settings for a Worker

To change the settings for a worker

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the server is located.
4. Click on the server.
5. Click **Workers**.
6. Go to the type of worker you have been directed to change.
7. If applicable, change Count or Memory Threshold using the dropdown boxes.
8. Click **Save**.

4.1.4.4. Database

Use the Database page to specify the location of your Virtual Machine Database (VMDB) and its login credentials. By default, the type is PostgreSQL on the Server.



NOTE

The server may not start if the database settings are changed. Be sure to validate your new settings before restarting the server.

4.1.4.4.1. Changing a Database Setting

To change a database setting:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the server is located.
4. Click on the server.
5. Click the **Database** tab.
6. In the **Database** area, select the **Type** of database. You can select from **External Database on another CFME appliance**, **External Postgres Database**, and **Internal Database on this CFME Appliance**.

- Use **Hostname** to specify the IP address or hostname of the external database server.
 - Use **Database Name** to specify the name of your VMDB.
 - Specify the **User Name** to connect to the VMDB.
 - Use **Password** and **Verify Password** to specify the password for the user name.
7. Click **Validate** to check the settings.
 8. Click **Save**.
 9. Click **OK** to the warning that the server will restart immediately after you save the changes.

During the restart, you are unable to access the server. When the restart is complete, the new database settings are in effect.

4.1.4.5. Customization and Logos

4.1.4.5.1. Custom Logos

Use **Custom Logos** to display your own logo in the corner of the CloudForms user interface and on the login screen. Use the procedures below to upload a custom logo to the user interface, and to customize the login background and login panel text on the user interface.



NOTE

- If you have upgraded from an earlier Red Hat CloudForms version and your custom logo was already in use before migration, although your logo image file is still in place in `vmdb/public/upload` you may have to uncheck and recheck the option to **Use Custom Logo Image** to re-enable displaying your custom logo. See [Section 4.1.4.5.2, “Uploading a Custom Logo to the User Interface”](#) for the procedure on how to access the **Use Custom Logo Image** option, or if you want to upload another custom logo to the user interface and customize the login background image and login panel text.
- Additionally, ensure the option to use configuration settings for the tenant under **Access Control** is set to **Yes**; see [Section 4.1.4.5.5, “Displaying the Custom Configuration Settings”](#) for the procedure on how to set the configuration settings.

4.1.4.5.2. Uploading a Custom Logo to the User Interface

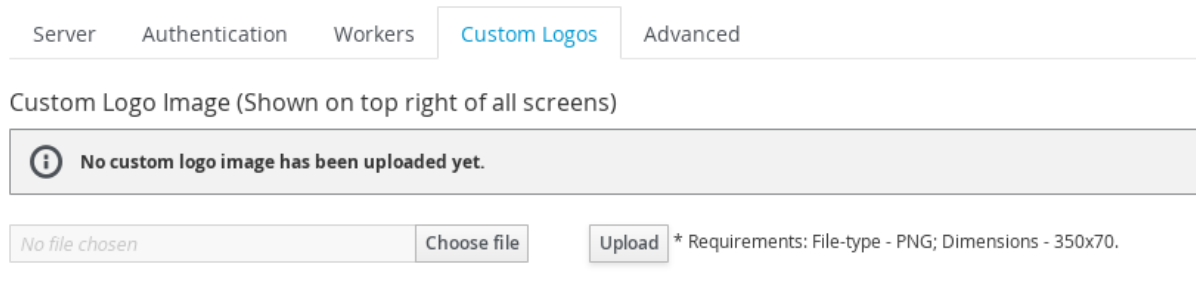


NOTE


Make sure the desired logo is accessible from the computer where you are running the CloudForms user interface. The file must be in portable network graphics (png) format with dimensions of 350 px x 70 px.

To upload a custom logo to the user interface:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the Red Hat CloudForms server is located.
4. Click on the server.
5. Click the **Custom Logos** tab.



6. In **Custom Logo Image (Shown on top right of all screens)**, click **Choose file** to go to the location where the logo file is located.
7. Click **Upload**. The icon is displayed above the file name box, and an option is shown to use the logo.
8. Check **Use Custom Logo Image** to add the logo to your user interface.
9. Click **Save**.



NOTE

To enable displaying your custom logo, ensure the option to use configuration settings for the tenant under **Access Control** is set to **Yes**. See [Section 4.1.4.5.5, “Displaying the Custom Configuration Settings”](#) for the procedure on how to set the configuration settings.

4.1.4.5.3. Customizing the Login Background



NOTE

Make sure the background image that you want to use is accessible from the computer where you are running the user interface. The file must be in PNG format with dimensions of 1280 px x 1000 px.

To customize the login background:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the server is located.
4. Click on the server.
5. Click the **Custom Logos** tab.

6. In **Custom Login & About Screen Background Image**, click **Choose file** to go to the location where the background image file is located.

Custom Login & 'About' Screen Background Image

i No custom login image has been uploaded yet.

No file chosen

Choose file

Upload

* Requirements: File-type - PNG; Dimensions - 1280x1000.

7. Click **Upload**. The icon is displayed above the file name box, and an option is shown to use the logo.
8. Check **Use Custom Login Background Image** to add the background image to the login screen of the user interface.
9. Click **Save**.

4.1.4.5.4. Customizing the Login Panel Text

To customize the login panel text:

1. From the settings menu, select **Configuration**.
2. Navigate to **Settings → Configuration**.
3. Click on the **Settings** accordion, then click **Zones**.
4. Click the zone where the server is located.
5. Click on the server.
6. Click the **Custom Logos** tab.
7. In **Custom Login Panel Text**, enter the text that you want to display on the login screen.
8. Click **Use Custom Login Text** to switch it to **Yes**.

Custom Login Panel Text (13 / 500)

Custom Text

Use Custom Login Text



☒ Yes

9. Click **Save**.

4.1.4.5.5. Displaying the Custom Configuration Settings

To enable displaying your custom logo in the corner of the Red Hat CloudForms user interface and on the login screen:

1. From the settings menu, select **Configuration**.
2. Click the **Access Control** accordion.
3. Click **Tenants**, then click **My Company**.

4. Click  (**Configuration**), then click  (**Edit this item**).

Editing Tenant "My Company"

Name	<input type="text" value="My Company"/>
Description	<input type="text" value="Tenant for My Company"/>
Use Configuration Settings	<input type="checkbox"/> No

5. Click **Use Configuration Settings** to switch it to **Yes**.
6. Click **Save**.

4.1.4.6. Advanced Settings

You may be instructed by Red Hat to edit some configuration settings manually. This feature is available for a limited number of options and can only be used by users assigned the super administrator role. Changing settings using this procedure may disable your Red Hat CloudForms server.



NOTE

Only make manual changes to your configuration files if directed to do so by Red Hat.

4.1.4.6.1. Editing Configuration Files Manually

To edit configuration files manually:

1. From the settings menu, click **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the server is located.
4. Click on the server.
5. Click the **Advanced** tab.
6. Select the configuration file to edit from the **Configuration File to Edit** area.
7. Make the required changes.
8. Click **Save**.

4.1.4.6.2. Configuration Parameters

Table: authentication

Parameters	Description
------------	-------------

Parameters	Description
amazon_key	If using Amazon for the authentication mode, specify your Amazon Key. This is the same as Amazon Access Key in Configuration-Operations-Server-Amazon Settings in the appliance console. Default: blank
amazon_secret	If using Amazon for the authentication mode, specify your Amazon Secret. This is the same as Amazon Secret Key in Configuration-Operations-Server-Amazon Settings in the appliance console. Default: blank
basedn	If using ldap for the authentication mode, specify your Base DN. This is the same as Base DN in Configuration-Operations- Server-LDAP Settings in the appliance console. Default: blank
bind_dn	The user name to bind to the LDAP server. This user must have read access to all users and groups that will be used for Red Hat CloudForms authentication and role assignment. This is the same as Bind DN in Configuration-Operations-Server-LDAP Settings in the appliance console. Default: blank
bind_pwd:	The password for the bind_dn user. This is the same as Bind Password in Configuration-Operations- Server-LDAP Settings in the appliance console. Default: blank
get_direct_groups	Use this to get the LDAP roles from the LDAP users' home forest. This is the same as Get Roles from Home Forest in the Authentication page for the Red Hat CloudForms Server. Default: true
group_memberships_max_depth	When traversing group memberships in the LDAP directory it will stop at this value. Default: 2
ldaphost	Use ldaphost to specify the fully qualified domain name of your LDAP server. This is the same as LDAP Host Name in Configuration-Operations-Server-LDAP Settings in the appliance console. Default: blank
ldapport	Specify the port of your LDAP server. This is the same as LDAP Port in Configuration-Operations- Server-LDAP Settings in the appliance console. Default: 389
mode	Use database to use the VMDB for security. Use ldap or ldaps to use directory services. This is the same as Mode in Configuration-Operations-Server-Authentication in the appliance console. Default: database

Parameters	Description
user_type	Use userprincipalname to type the user name in the format of user@domainname. Use mail to login with the user's e-mail address. Use dn-cn for Distinguished Name (CN=<user>) or dn-uid Distinguished Name (UID=<user>) to use just the user name, but be sure to enter the proper user_suffix for either one. This is the same as User Type in Configuration-Operations- Server-LDAP Settings in the appliance console. Default: userprincipalname
user_suffix	Domain name to be used with user_type of dn-cn or dn-uid. This is the same as User Suffix in Configuration-Operations- Server-LDAP Settings in the appliance console. Default: blank

Table: coresident_miqproxy

Parameters	Description
use_vim_broker	Specify if you want the coresident SmartProxy to use a shared connection through the VIM broker to communicate with the VC or ESX host for SmartState Analysis. If it is disabled, then each SmartProxy SmartState Analysis would create its own connection. Default: true
concurrent_per_ems	Specify the number of co-resident SmartProxy SmartState Analyses that can be run against a specific management system at the same time. Default: 1
concurrent_per_host	Specify the number of co-resident SmartProxy SmartState Analyses that can be run against a specific host at the same time. Default: 1
scan_via_host	If you change scan_via_host to false, Red Hat CloudForms will use the Management System to scan which is limited by the concurrent_per_ems setting instead of the concurrent_per_host setting. Note this will greatly increase traffic to the Management System. Default: true

Table: ems_refresh

Parameters	Description
------------	-------------

Parameters	Description
capture_vm_created_on_date	Set to false to turn off historical event retrieval. Set to true to turn on. By setting the flag to true Red Hat CloudForms will try to set the "ems_created_on" column in the vms table after an ems refresh for new VMs and any VMs with a nil "ems_created_on" value. Red Hat CloudForms looks at event information in our database as well as looking up historical event data from the management system. This is optional since the historical lookup could timeout. Default: false
collect_advanced_settings	Set to false if you do not want to collect advanced Virtual Machine settings during a management system refresh. This will increase the speed of the refresh, but less data will be collected. If the parameter is not listed, then the value is true. Default: true
ec2	
get_private_images	For EC2 refreshes only; whether or not to retrieve private images. Default: true
get_public_images	For EC2 refreshes only; whether or not to retrieve public images. Default: false. Warning: setting get_public_images to true loads several thousand images in the VMDB by default and may cause performance issues.
get_shared_images	For EC2 refreshes only; whether or not to retrieve shared images. Default: true
public_images_filters	For EC2 refreshes only; a filter to reduce the number of public images. Default: all public images
ignore_terminated_instances	For EC2 refreshes only; whether or not to ignore terminated instances. Default: true
full_refresh_threshold	The number of targeted refreshes requested before they are rolled into a full refresh. For example, if the system and/or the user target a refresh against 7 VMs and 2 Hosts (9 targets), when the refresh actually occurs it will do a partial refresh against those 9 targets only. However, if a 10th had been added, the system would perform a full EMS refresh instead. Default: 100
raise_vm_snapshot_complete_if_created_within:	Raises vm_snapshot_complete event for a snapshot being added to VMDB only if the create time in Virtual Center is within the configured period of time. This prevents raising events for old snapshots when a new VC is added to Red Hat CloudForms. Default: 15.minutes
refresh_interval	Scheduler does a periodic full EMS refresh every refresh_interval. Default: 24.hours

Table: host_scan

Parameters	Description
queue_timeout	Time period after which a host SmartState analysis will be considered timed out. Default: 20.minutes

Table: log

Parameters	Description
level	Specify the required level of logging for the Red Hat CloudForms appliance. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. This is the same as Log Level in Configuration-Operations-Server-Logging in the appliance console and applies immediately to the evm.log file. Default: info
level_aws	Specify the level of logging for Amazon Web Services communications. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. This applies to the aws.log file. Default: info
level_aws_in_evm	Specify what level of Amazon Web Services communication log should be also shown in evm.log. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. Default: error
level_fog	Specify the level of logging for Fog communications. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. This applies to the fog.log file. Default: info
level_fog_in_evm	Specify what level of Fog communication log should be also shown in evm.log. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. Default: error
level_rails	Specify the level of logging for Rails. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. Once changed, this applies immediately to the production.log file. Default: info
level_rhevm	Specify the level of logging for Red Hat communications. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. This applies to the rhevm.log file. Default: warn
level_rhevm_in_evm	Specify what level of Red Hat communication log should be also shown in evm.log. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. Default: error

Parameters	Description
level_vim	Specify the level of logging for VIM (communication with VMware ESX and Virtual Center). Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. This applies to the vim.log file. Default: warn
level_vim_in_evm	Specify what level of vim logging should be also shown in evm.log. Possible levels from most detailed to least detailed are: debug, info, warn, error, fatal. Default: error

Table: db_stats

Parameters	Description
enabled	Specify if you want to keep track of the number of queries, size of queries, number of responses, size of response, min/max for each, number of established connections at for each server process. This information will show in the EVM log. Default: false
log_frequency	How frequently in seconds the process will log the database statistic in seconds. Default: 60

Table 3.7. callsites

Table: log

Parameters	Description
enabled	Specify if you want keep track of the code that is accessing the database. Enabling call sites will decrease performance because of the amount of information tracked. The db_stats: enabled parameter must be set to true to use this. Default: false
depth	Specify how many levels in the call stack to track for each database access. Default: 10
min_threshold	Do not keep track of code that does not access the database this many times per log_frequency. Default: 10
path	Set the path for the Red Hat CloudForms appliance log. This is the same as Log Path in Configuration-Operations- Server-Logging in the appliance console. Default: If no value is present, the path is /var/www/miq/vmdb/log.
line_limit	Limit how many characters are retained in a single log line. 0 means no limit. Default: 0

Table 3.8. collection

Parameters	Description
ping_depot	Whether to use TCP port ping to the log depot before performing log collection. Default: true
ping_depot_timeout	Specify how long in seconds to wait for response from log depot before deciding that the TCP port ping failed. Default: 20
current	<p>When collecting logs, specifies what is considered current logging as opposed to archived logging. Default: :pattern:</p> <p>log/*.log</p> <p>log/apache/*.log</p> <p>log/*.txt</p> <p>config/*</p> <p>/var/opt/rh/rh-postgresql94/lib/pgsql/data/*.conf</p> <p>/var/opt/rh/rh-postgresql94/lib/pgsql/data/pg_log/*</p> <p>/var/log/syslog*</p> <p>/var/log/daemon.log*</p> <p>/etc/default/ntp*</p> <p>/var/log/messages*</p> <p>/var/log/cron*</p> <p>BUILD</p> <p>GUID</p> <p>VERSION</p>
archive	Specifies what is considered archived logging. The default pattern is blank which means *.gz files in the log directory.

Table 3.9. log_depot

Parameters	Description
uri	Specify the URI for the log depot. This is the same as the URI in Configuration → Diagnostics → Collect Logs in the appliance console. Default: blank

Parameters	Description
username	Specify the user name for the log depot. This is the same as the user ID in Configuration → Diagnostics → Collect Logs in the appliance console. Default: blank
password	Specify the password for the user for the log depot. This is the same as the password in Configuration → Diagnostics → Collect Logs in the appliance console. Default: blank

Table: performance

Parameters	Description
capture_threshold	
vm	Amount of time in minutes to wait after capture before capturing again. Default: 50.minutes
host	Amount of time in minutes to wait after capture before capturing again. Default: 50.minutes
ems_cluster	Amount of time in minutes to wait after capture before capturing again. Default: 50.minutes
storage	Amount of time in minutes to wait after capture before capturing again. Default: 120.minutes
capture_threshold_with_alerts	
host	Amount of time in minutes to wait after capture before capturing again. This value is used instead of capture_threshold for Hosts that have alerts assigned based on real time Capacity & Utilization data. Default: 20.minutes
ems_cluster	Amount of time in minutes to wait after capture before capturing again. This value is used instead of capture_threshold for clusters that have alerts assigned based on real time Capacity & Utilization data. Default: 50.minutes
vm	Amount of time in minutes to wait after capture before capturing again. This value is used instead of capture_threshold for VMs that have alerts assigned based on real time Capacity & Utilization data. Default: 20.minutes
concurrent_requests	

Parameters	Description
vm	Amount of time in minutes to wait after capture before capturing again. This value is used instead of capture_threshold for VMs that have alerts assigned based on real time Capacity & Utilization data. Default: 20.minutes
hourly	Number of concurrent VC requests to make when capturing hourly raw metrics. Default: 1
realtime	Number of concurrent VC requests to make when capturing real time raw metrics. Default: 20
history	
initial_capture_days	How many days to collect data for on first collection. Default: 0
Keep_daily_performances	How long to keep daily performance data in the VMDB. Default: 6.months
keep_realtime_performances	How long to keep realtime performance data in the VMDB. Default: 4.hours
keep_hourly_performances	How long to keep hourly performance data in the VMDB. Default: 6.months
purge_window_size	When the purge needs to delete rows which are older than the keep_realtime_performances, keep_hourly_performances, and keep_daily_performances values, this value sets how many rows to delete in each batch. For example, a value of 1000 will cause us to issue ten 1,000 row deletes. Default: 1000

Table 3.11. repository_scanning

Parameters	Description
defaultsmartproxy	Specify the SmartProxy for repository scanning. This is the same as Default Repository Smartproxy in Configuration-Operations-Server-VM Server Control in the appliance console. Default: blank

Table 3.12. server

Parameters	Description
case_sensitive_name_search	Specify if you want the search by name on configuration item screens to be case sensitive. Default: false

Parameters	Description
company	Specify the label you want to use for your company's tagging. This is the same as Company Name in Configuration-Operations-Server-Basic Info. Default: "My Company"
custom_logo	Specify if you want to use a custom logo. This is the same as Use Custom Logo in Configuration-Custom Logo-Logo Selection. Default: false
events	
disk_usage_gt_percent	For Red Hat CloudForms operational alerts, specify at what threshold the disk usage alerts will be triggered. Default: 80
heartbeat_timeout	How long to wait until the server heartbeat is considered timed out. if the timeout is exceeded, other appliances in the zone/region can vie for the roles active on the timed out Red Hat CloudForms appliance. Default: 2.minutes
host	Red Hat CloudForms Server's IP address. Default: blank
hostname	Red Hat CloudForms Server's hostname. Default: localhost.localdomain
listening_port	Specify the port number on which the web server is listening. Note that this does not set the port that VMDB listens on. When deploying the SmartHost from the Red Hat CloudForms appliance, it tells the SmartHost (miqhost) what port to talk to the VMDB on. Default: "443"
mks_version	Specify the version of the VMware MKS Plugin to use for the VM Console. This is the same as VMware MKS Plugin Version in Configuration-Operations- Server-VM Console. Default : 2.1.0.0
name	Set the name to display for the Red Hat CloudForms appliance that you are logged on to in the appliance console. This is the same as appliance Name in Configuration-Operations- Server-Basic Information. Default : EVM

Parameters	Description
role	Specify the roles for this Red Hat CloudForms Server, separated by commas without spaces. The possible values are automate, database_operations, ems_inventory, ems_metrics_collector, ems_metrics_coordinator, ems_metrics_processor, ems_operations, event, notifier, reporting, scheduler, smartproxy, smartstate, user_interface, web_services. This is the same as Server Roles in Configuration-Operations- Server- Server Control. Default: database_operations, event, reporting, scheduler, smartstate, ems_operations, ems_inventory, user_interface, web_services
session_store	Where to store the session information for all web requests. The possible values are sql, memory, or cache. SQL stores the session information in the database regardless of the type of database server. Memory stores all the session information in memory of the server process. Cache stores the information in a memcache server. Default: cache
startup_timeout	The amount of time in seconds that the server will wait and prevent logins during server startup before assuming the server has timed out starting and will redirect the user to the log page after login. Default: 300
timezone	Set the timezone for the Red Hat CloudForms appliance. Default: UTC
vnc_port	If using VNC for remote console, the port used by VNC. Default: 5800
zone	Set the Zone for this appliance belongs. This is the same as Zone in Configuration-Operations- Server-Basic Information. Default : default
:worker_monitor	Starts and monitors the workers. Parameters specified here will override those set in the workers:default section.
poll	How often the worker monitor checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 15.seconds
miq_server_time_threshold	How much time to give the server to heartbeat before worker monitor starts to take action against non-responding server. Default: 2.minutes

Parameters	Description
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 1
sync_interval	Time interval to sync active roles and configuration for all workers. Default: 30.minutes
wait_for_started_timeout	How long to wait for a started worker to heartbeat before considering the worker timed out. Default: 10.minutes
kill_algorithm	
name	Criteria used to start killing workers. Default: used_swap_percent_gt_value
value	Value of the criteria used. Default: 80
start_algorithm	
name	After server startup, criteria that must be met to decide if the Red Hat CloudForms Server can start a new worker. Default: used_swap_percent_lt_value
value	Value of criteria used. Default: 60

Table: session

Parameters	Description
interval	Set the time interval in seconds for checking inactive sessions in appliance console. Default: 60
timeout	Set the time period in seconds in which inactive console sessions are deleted. This is the same as Session Timeout in Configuration-Operations-Server-Authentication in the appliance console. Default: 3600
memcache_server	If you choose memory for session_store, you need to specify the memcache_server to retrieve the session information from. Default: 127.0.1.1:11211
memcache_server_opts	Options to send to memcache server. : blank
show_login_info	Specify whether or not you want to see login info on start page. Default: true

Table: smartproxy_deploy

Parameters	Description
queue_timeout	Timeout for host smartproxy deploy job. Default: 30.minutes

Table 3.15. smtp

Parameters	Description
host	Specify the hostname of the smtp mail server. This is the same as Host in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: localhost
port	Specify the port of the smtp mail server. This is the same as Port in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: "25"
domain	Specify the domain of the smtp mail server. This is the same as Domain in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: mydomain.com
authentication	Specify the type of authentication of the smtp mail server. This is the same as Authentication in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: login
user_name	Specify the username required for login to the smtp mail server. This is the same as User Name in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: evmadmin
password	Specify the encrypted password for the user_name account. This is the same as Password in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: blank
from	Set the address that you want to send e-mails from. This is the same as From E-mail Address in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: cfadmin@cfserver.com

Table 3.16. snapshots

Parameters	Description
create_free_percent	Ensures the % of free space available on the main datastore (datastore where vmx file is located) can support the % growth of the snapshot. The default is to require space for 100% of the provisioned size of all disks that are taking part in the snapshot. A value of 0 means do not check for space before creating the snapshot. Default: 100

Parameters	Description
remove_free_percent	Ensures the % of free space available on the main datastore (datastore where vmx file is located) has the % free space available to support the snapshot deletion process. Note that the deletion process consists of first composing a new snapshot then removing it once the original snapshot to be deleted has been collapsed in the VM. The default is to require 100% of the size of all disks to complete this process. A value of 0 means do not check for space before removing the snapshot. Default: 100

Table 3.17. webservices

Parameters	Description
contactwith	Set to ipaddress to contact miqhost using the IP address. Set to hostname to contact miqhost by its hostname. Set to resolved_ipaddress to take the hostname and resolve it to an IP address. Default: ipaddress
mode	Set to invoke to use webservices. Set to disable to turn off webservices. This is the same as Mode in Configuration-Operations- Server-Web Services in the appliance console. Default: invoke
nameresolution	If set to true, the hostname will be resolved to an IP address and saved with the host information in the VMDB. Default: false
security	If Web Services are enabled, you can set this to ws-security. This is the same as Security in Configuration-Operations- Server-Web Services in the appliance console. Note: This is not currently supported. Default: none
timeout	Specify the web service timeout in seconds. Default: 120
password	Specify the encrypted password for the user_name account. This is the same as Password in Configuration-Operations-Server-Outgoing SMTP E-mail Server. Default: blank
use_vim_broker	Controls if the vim_broker is used to communicate with VMware infrastructure. Default: true

Table: workers

Parameters	Description
worker_base	
defaults	If the following parameters are NOT explicitly defined for a specific worker, then these values will be used.
count	Number of this type of worker. Default: 1
gc_interval	How often to do garbage collection for this worker. Default: 15.minutes
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 3.seconds
poll_method	If set to normal, the worker checks for work the number of seconds set in the poll parameter. If set to escalate, the worker will increase the time between checks when there is no work to be done. Default: normal
poll_escalate_max	The maximum number of time to wait between checks for work. Poll_method must be set to escalate for this option to be used. Default: 30.seconds
heartbeat_freq	How often to "heartbeat" the worker. Default: 60.seconds
heartbeat_method	Set which way to dispatch work. Possible values are sql or drb. Default: drb
heartbeat_timeout	How long to wait until the worker heartbeat is considered timed out. Default: 2.minutes
parent_time_threshold	How long to allow the parent to go without heartbeating before considering the "parent" not responding. For workers, the worker monitor is the parent. For Worker monitor, the Server is the parent. Default: 3.minutes
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 150.megabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 10

Parameters	Description
restart_interval	How long to let a worker remain up before asking it to restart. All queue based workers are set to 2.hours and every other worker does not get restarted by a 0.hours value. Default: 0.hours
starting_timeout	How long to wait before checking a worker's heartbeat when it is starting up to mark it as not responding, similar to a grace period before you begin monitoring it. Default: 10.minutes
event_catcher	Associated with Event Monitor Server Role. Captures ems events and queues them up for the event_handler to process. Parameters specified here will override those set in the worker_base:default section.
ems_event_page_size	Internal system setting which sets the maximum page size for the event collector history. This should not be modified. Default: 100
ems_event_thread_shutdown_timeout	Internal system setting which determines how long the event catcher at shutdown will wait for the event monitor thread to stop. This should not be modified. Default: 10.seconds
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 2.gigabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 1
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 1.seconds
event_catcher_redhat	Contains settings that supersede the event_catcher for event_catcher_redhat.
event_catcher_vmware	Contains settings that supersede the event_catcher for event_catcher_vmware.
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 1.seconds
event_catcher_openstack	Contains settings that supersede the event_catcher for event_catcher_openstack.

Parameters	Description
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 15.seconds
topics	List of AMQP topics that should be monitored by Red Hat CloudForms when gathering events from OpenStack.
duration	Qpid Specific. Length of time (in seconds) the receiver should wait for a message from the Qpid broker before timing out. Default: 10.seconds
capacity	Qpid Specific. The total number of messages that can be held locally by the Qpid client before it needs to fetch more messages from the broker. Default: 50.seconds
amqp_port	Port used for AMQP. Default: 5672
schedule_worker	Settings for Scheduler Server Role and any other work that runs on a schedule. Parameters specified here will override those set in the worker_base:default section.
db_diagnostics_interval	How frequently to collect database diagnostics statistics. Default: 30.minutes
job_proxy_dispatcher_interval	How often to check for available SmartProxies for SmartState Analysis jobs. Default: 15.seconds
job_proxy_dispatcher_stale_message_check_interval	How often to check for the dispatch message in the queue Default: 60.seconds
job_proxy_dispatcher_stale_message_timeout	Kill a message if this value is reached. Default: 2.minutes
job_timeout_interval	How often to check to see if a job has timed out. Default: 60.seconds
license_check_interval	How often to check for valid license. Default: 1.days
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 150.megabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 3

Parameters	Description
performance_collection_interval	Controls how often the schedule worker will put performance collection request on the queue to be picked up by the collection worker. Default: 3.minutes
performance_collection_start_delay	How long after Red Hat CloudForms Server has started before starting capacity and utilization collection, if collection needs to be done. Default: 5.minutes
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 15.seconds
server_logs_stats_interval	How often to log the Red Hat CloudForms Server statistics. Default: 5.minutes
server_stats_interval	How often to collect the Red Hat CloudForms Server statistics. Default: 60.seconds
session_timeout_interval	How often to check to see if a UI (appliance console) session has timed out. Default: 30.seconds
storage_file_collection_interval	How often to perform file inventory of storage locations. Default: 1.days
storage_file_collection_time_utc	What time to perform file inventory of storage locations. Default: "06:00"
vdi_refresh_interval	How often to refresh vdi inventory. Default: 20.minutes
vm_retired_interval	How often to check for virtual machines that should be retired. Default: 10.minutes
vm_scan_interval	How often to check virtual machines to see if scan needs to be done. Default: 10.minutes
smis_refresh_worker	Settings for Storage Inventory Server Role and any other work that runs on a schedule. Parameters specified here will override those set in the worker_base:default section
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 15.seconds

Parameters	Description
connection_pool_size	Maximum number of database connections allowed per process. Default: 5
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 1.gigabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 3
smis_update_period	How frequently to update smis information. Default: 1.hours
status_update_period	How frequently to update smis status. Default: 5.minutes
stats_update_period	How frequently to update smis statistics. Default: 10.minutes
vim_broker_worker	Launched for any of these roles: Capacity & Utilization Collector, SmartProxy, SmartState Analysis, Management System Operations, Management System Inventory. Also launched if the use_vim_broker setting is on. Provides connection pooling, caching of data to and from the VMware infrastructure. Parameters specified here will override those set in the workers:default section.
heartbeat_freq	How often to heartbeat the worker. Default: 15.seconds
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 1.gigabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 3
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 1.seconds
reconnect_retry_interval	Period after which connection is retried. Default: 5.minutes
vim_broker_status_interval	Internal system setting which configures how much time to wait after receiving event updates before checking for more updates. Default: 0.seconds
wait_for_started_timeout	Time between the worker's preload and startup time before considering the worker timed out. Default: 10.minutes

Parameters	Description
ui_worker:	Settings for User Interface Server Role. Parameters specified here will override those set in the worker_base:default section.
connection_pool_size	Maximum number of database connections allowed per process. Default: 5
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 1.gigabytes
nice_delta: 1	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 1
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 60.seconds
web_service_worker	Settings for Web Services Server Role. Parameters specified here will override those set in the worker_base:default section.
connection_pool_size	Maximum number of database connections allowed per process. Default: 5
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 1.gigabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 1
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 60.seconds
queue_worker_base	Base class of all queue workers that work off of the queue..
defaults	If the following parameters are NOT explicitly defined for a queue worker, then these values will be used.
cpu_usage_threshold	How much cpu to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 100.percent
queue_timeout	How long a queue message can be worked on before it is considered timed out. Default: 10.minutes

Parameters	Description
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 400.megabytes
restart_interval	Queue workers restart interval. Default: 2.hours
poll_method	If set to normal, the worker checks for work the number of seconds set in the poll parameter. If set to escalate, the worker will increase the time between checks when there is no work to be done. Default: normal
generic_worker	Performs work that is not classified as any specific type of work. Processes all normal priority or non-specific queue items. Parameters specified here will override those set in the queue_worker_base:default section
count	Number of this type of worker. Default: 4
ems_refresh_worker	Performs all ems (management system) refreshes to keep the vmdb in sync with the state of the components of the virtual infrastructure in the various management systems. Parameters specified here will override those set in the queue_worker_base:default section
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 10.seconds
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 2.gigabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 7
restart_interval	Queue workers restart interval. Default: 2.hours
queue_timeout	How long a message can be worked on before it is considered timed out. Default: 120.minutes
event_handler	Associated with Event Monitor Server Role. Handles all events caught by the event catcher worker. Parameters specified here will override those set in the workers:default section. Parameters specified here will override those set in the queue_worker_base:default section

Parameters	Description
cpu_usage_threshold	How much cpu to allow the worker to grow to before gracefully requesting it to exit and restart. The value of 0 means that this worker will never be killed due to CPU usage. Default: 0.percent
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 7
perf_collector_worker	Connects to VC/ESX to collect the raw performance data. Same as the Capacity & Utilization Data Collector Server Role. Parameters specified here will override those set in the queue_worker_base:default section count. Number of this type of worker. Default: 2
count	Number of this type of worker. Default: 2
poll_method	If set to normal, the worker checks for work the number of seconds set in the poll parameter. If set to escalate, the worker will increase the time between checks when there is no work to be done. Default: escalate
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 3
perf_processor_worker	Processes the raw performance data into a reportable format. Same as the Capacity & Utilization Data Processor Server Role. Parameters specified here will override those set in the queue_worker_base:default section
count	Number of this type of worker. Default: 2
poll_method	If set to normal, the worker checks for work the number of seconds set in the poll parameter. If set to escalate, the worker will increase the time between checks when there is no work to be done. Default: escalate
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 400.megabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 7
priority_worker	Performs all high priority queue items including many tasks on behalf of the UI. UI requests are normally executed by a priority worker so they will not to block the UI. Parameters specified here will override those set in the queue_worker_base:default section

Parameters	Description
count	Number of this type of worker. Default: 2
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 200.megabytes
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 1
poll	How often the workers checks for work. This value only is only used when the worker has no more work to do from the queue. It will wait for an amount of time determined by the poll value and poll method. Therefore, if there is constant work on the queue, the worker will not wait in between messages. Default: 1.seconds
reporting_worker	Compiles reports. Settings for Reporting Server Role. Parameters specified here will override those set in the queue_worker_base:default section
count	Number of this type of worker. Default: 2
nice_delta	Tells the worker monitor what Unix "nice" value to assign the workers when starting. A lower number is less nice to other processes. Default: 7
smart_proxy_worker	Performs the embedded scanning of virtual machines. Settings for SmartProxy Server Role. Parameters specified here will override those set in the queue_worker_base:default section
count	Number of this type of worker. Default: 3
memory_threshold	How much memory to allow the worker to grow to before gracefully requesting it to exit and restart. Default: 600.megabytes
queue_timeout	How long a queue message can be worked on before it is considered timed out. Default: 20.minutes
restart_interval	Queue workers restart interval. Default: 2.hours



4.1.5. Schedules

4.1.5.1. Scheduling SmartState Analyses and Backups

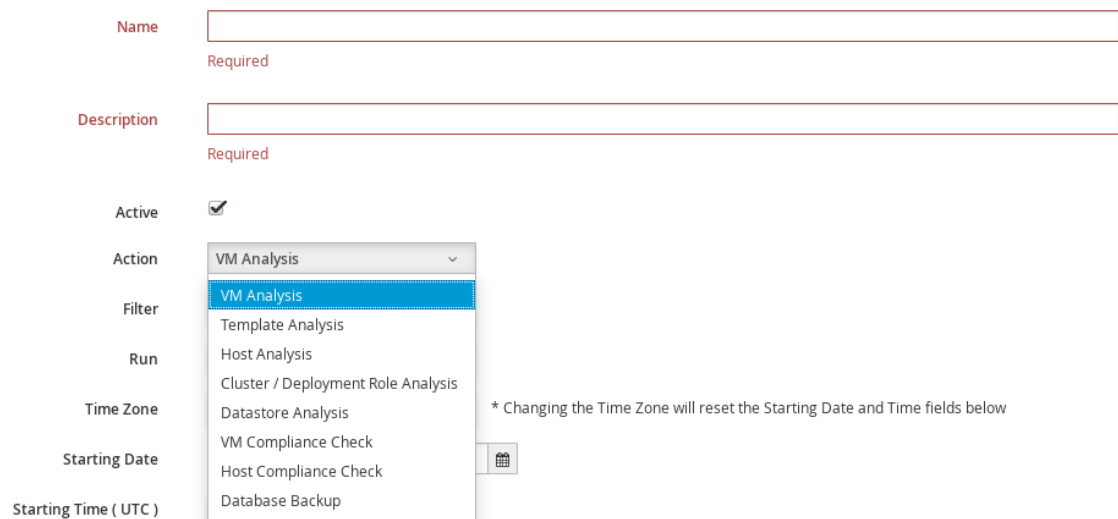
From the **Schedules** area in **Settings**, you can schedule the analyses of virtual machines, hosts, clusters, and datastores to keep the information current. Depending on which resource you want to analyze, you can filter which ones to analyze. You may also specify only one virtual machine or perform an analysis on all virtual machines. In addition, you can schedule compliance checks, and database backups.

4.1.5.1.1. Scheduling a SmartState Analysis or Compliance Check

To schedule a SmartState Analysis or Compliance Check:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Schedules**.
3. Click  (**Configuration**), and  (**Add a new Schedule**).
4. In the **Basic Information** area, type in a **Name** and **Description** for the schedule.
5. Select **Active** to enable this scan.
6. From the **Action** list, select the type of analysis to schedule. Based on the type of analysis you choose, you are presented with one of the following group boxes:

Adding a new Schedule



The screenshot shows a form titled 'Adding a new Schedule'. It has several fields: 'Name' (required), 'Description' (required), 'Active' (checkbox, checked), 'Action' (dropdown menu), 'Filter', 'Run', 'Time Zone', 'Starting Date', and 'Starting Time (UTC)'. The 'Action' dropdown is open, showing a list of options: 'VM Analysis', 'Template Analysis', 'Host Analysis', 'Cluster / Deployment Role Analysis', 'Datastore Analysis', 'VM Compliance Check', 'Host Compliance Check', and 'Database Backup'. A note at the bottom right of the dropdown says '* Changing the Time Zone will reset the Starting Date and Time fields below'.

- **VM Analysis:** Displays **VM Selection** where you can choose to analyze **All VMs**, **All VMs for Provider**, **All VMs for Cluster**, **All VMs for Host**, **A single VM**, or **Global Filters**.
- **Template Analysis:** Displays **Template Selection** where you can choose to analyze **All Templates**, **All Templates for Provider**, **All Templates for Cluster**, **All Templates for Host**, **A single Template**, or **Global Filters**.
- **Host Analysis:** Displays **Host Selection** where you can choose to analyze **All Hosts**, **All Hosts for Provider**, **A single Host**, or **Global Filters**.



NOTE

You can only schedule host analyses for connected virtual machines, not repository virtual machines that were discovered through that host. Since repository virtual machines do not retain a relationship with the host that discovered them, there is no current way to scan them through the scheduling feature. The host is shown because it may have connected virtual machines in the future when the schedule is set to run.

- **Cluster / Deployment Role Analysis:** Displays **Cluster Selection** where you can choose to analyze **All Clusters**, **All Clusters for Provider**, or **A single Cluster**.
 - **Datastore Analysis:** Displays **Datastore Selection** where you can choose to analyze **All Datastores**, **All Datastores for Host**, **All Datastores for Provider**, **A single Datastore**, or **Global Filters**.
 - **VM Compliance Check:** Displays **VM Selection** where you can choose to analyze **All VMs**, **All VMs for Provider**, **All VMs for Cluster**, **All VMs for Host**, **A single VM**, or **Global Filters**.
 - **Host Compliance Check:** Displays **Host Selection** where you can choose to analyze **All Hosts**, **All Hosts for Provider**, **All Hosts for Cluster**, **A single Host**, or **Global Filters**.
 - **Database Backup:** Under **Type**, displays **Network File System** and **Samba**. See [Section 4.1.5.2, “Scheduling a Database Backup”](#) for details on scheduling a database backup.
- By applying **Global Filters** within any of the above items, you can designate which virtual machines or hosts to analyze.
 - In **Run**, set the frequency of the analysis to run. There are further options based on which **Run** option you choose.
 - Click **Once** to have the analysis run only one time.
 - Click **Daily** to run the analysis on a daily basis. You will be prompted to select the number of days between each analysis.
 - Click **Hourly** to run the analysis hourly. You will be prompted to select the number of hours between each analysis.

- Select a **Time Zone**.





NOTE

If you change the **Time Zone**, you will need to reset the starting date and time.

- Type or select a date to begin the schedule in **Starting Date**.
- Select a **Starting Time** based on a 24 hour clock in the selected Time Zone.
- Click **Add**.

4.1.5.2. Scheduling a Database Backup

To schedule a database backup:

- From the settings menu, select **Configuration**.
- Click on the **Settings** accordion, then click **Schedules**.
- Click  (**Configuration**), and  (**Add a new Schedule**).

4. In the **Basic Information** area, type in a **Name** and **Description** for the schedule.

Adding a new Schedule

Name	<input type="text" value="DB daily backup"/>		
Description	<input type="text" value="DB daily backup"/>		
Active	<input checked="" type="checkbox"/>		
Action	<input type="text" value="Database Backup"/>		
Type	<input type="text" value="Network File System"/>		
Depot Name	<input type="text"/>		
	Required		
URI	<input type="text" value="nfs://"/>		
	Required		
Run	<input type="text" value="Daily"/>	every	<input type="text" value="Day"/>
Time Zone	<input type="text" value="(GMT+00:00) UTC"/> * Changing the Time Zone will reset the Starting Date and Time fields below		
Starting Date	<input type="text" value="04/12/2016"/>		
Starting Time (UTC)	<input type="text" value="0"/>	h	<input type="text" value="0"/> m

5. Select **Active** to enable this backup schedule.
6. From the **Action** list, select **Database backup**.
7. In the **Database Backup Settings** area, select a type of server to put the backups. You can either use **Network File System** or **Samba**.
 - If selecting **Samba**, enter the **Depot Name**, **URI**, **User ID**, and a valid **Password**. Then, click **Validate** to check the settings.
 - If you choose **Network File System**, enter the **Depot Name** and **URI**.
8. In **Run**, set the frequency of the analysis to run. There are further options based on which **Run** option you choose.
 - Click **Once** to have the backup run only one time.
 - Click **Daily** to run the backup on a daily basis. You will be prompted to select the number of days between each backup.
 - Click **Hourly** to run the backup hourly. You will be prompted to select the number of hours between each backup.
9. Select a **Time Zone**.





NOTE

If you change the **Time Zone**, you will need to reset the stating date and time.

10. Type or select a date to begin the schedule in **Starting Date**.
11. Select a **Starting Time** (UTC) based on a 24 hour clock in the selected time zone.
12. Click **Add**.


4.1.5.2.1. Modifying a Schedule

To modify a schedule:

1. From the settings menu, select **Configuration**.
2. Click on the **Settings** accordion, then click **Schedules**.
3. Click the schedule that you want to edit.
4. Click  (**Configuration**), and then click  (**Edit this Schedule**).
5. Make the required changes.
6. Click **Save**.

4.2. ACCESS CONTROL

From the settings menu, select **Configuration**. Click on the **Access Control** accordion to see a hierarchy of the configurable items for users, groups, roles, and tenants. You can add and modify users, groups, account roles, tenants, and projects.





NOTE

For information about tenancy in CloudForms, and the difference between a tenant and project, see [Tenancy](#) in the *Deployment Planning Guide*.

4.2.1. Creating a Tenant



To create a tenant:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Tenants**.
3. Click on the top-level **Tenant**, click  (**Configuration**), and  (**Add child Tenant to this Tenant**) to create a tenant.
4. Enter a name for the tenant in the **Name** field.
5. Enter a description for the tenant in the **Description** field.
6. Click **Add**.

4.2.2. Creating a Project

To create a project:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Tenants**.

- Click on the **Tenant** where you want to add a **Project**, click  (**Configuration**), and  (**Add Project to this Tenant**) to create a project.
- Enter a name for the project in the **Name** field.
- Enter a description for the project in the **Description** field.
- Click **Add**.



4.2.3. Managing Tenant and Project Quotas

Use the following procedure to allocate or edit quotas for tenants and projects.



NOTE

Quota is allocated based on the user's current group. If the user belongs to multiple groups, you must change to the desired group before allocating or editing group quota. See [Section 4.2.7, "Groups"](#) for more details.

- From the settings menu, select **Configuration**.
- Click on the **Access Control** accordion, then click **Tenants**.
- Click on the **Tenant** or **Project** where you want to add a quota, click  (**Configuration**), and  (**Manage quotas for the Selected Item**) to create a quota.
- In the list of pre-built quotas, switch **Enforced** next to the quota item you want to enable to **Yes**.
- In the **Value** field, enter the constraints you want to apply to the quota.

Manage quotas for Tenant "My Company"

Enforced	Description	Value	Units
<input checked="" type="checkbox"/>	Allocated Virtual CPUs	30	Count
<input checked="" type="checkbox"/>	Allocated Memory in GB	100	GB
<input checked="" type="checkbox"/>	Allocated Storage in GB	1000	GB
<input type="checkbox"/>	Allocated Number of Virtual Machines	Not enforced	Count
<input type="checkbox"/>	Allocated Number of Templates	Not enforced	Count

[Save](#) [Reset](#) [Cancel](#)

- Click **Save**.

4.2.4. Tagging Tenants and Projects



To tag tenants and projects:

- From the settings menu, select **Configuration**.
- Click on the **Access Control** accordion, then click **Tenants**.
- Select the tenant or project, then click **Policy**, and select **Edit My Company Tags for this Tenant**.

4. In **Tag Assignment**, click **Select a customer tag to assign**, and select a tag from the list. In the next column, set a corresponding value.
5. Click **Save**.

4.2.5. Creating a User

To create a user:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Users**.
3. Click  (**Configuration**), and  (**Add a new User**) to create a user.
4. Enter a **Full Name**, **Username**, **Password** with confirmation, **Email Address**, and choose a **Group** for the user. An email address is required; omitting the user email address can result in unsuccessful provisioning requests.

Adding a new User

User Information

Full Name	<input type="text"/>
Username	<input type="text"/>
Password	<input type="password"/>
Confirm Password	<input type="password"/>
E-mail Address	<input type="text"/>
Group	<input data-bbox="564 1234 898 1272" type="text" value=" <Choose a Group> "/>



IMPORTANT

- If you are using LDAP, but did not enable **Get User Groups from LDAP** in your server's **Authentication** tab, you will need to define a user. The UserID must match exactly the user's name as defined in your directory service. Use all lowercase characters to ensure the user can be found in the VMDB.
- When the user logs in, they use their LDAP password.



5. Click **Add**.

4.2.6. Deleting a User

For security reasons, delete any user that no longer needs access to the information or functions of the server.

To delete a user:

1. From the settings menu, select **Configuration**.

2. Click on the **Access Control** accordion, then click **Users**.
3. Select the user you want to delete.
4. Click  (**Configuration**), and  (**Delete selected Users**) to delete a user.

4.2.7. Groups

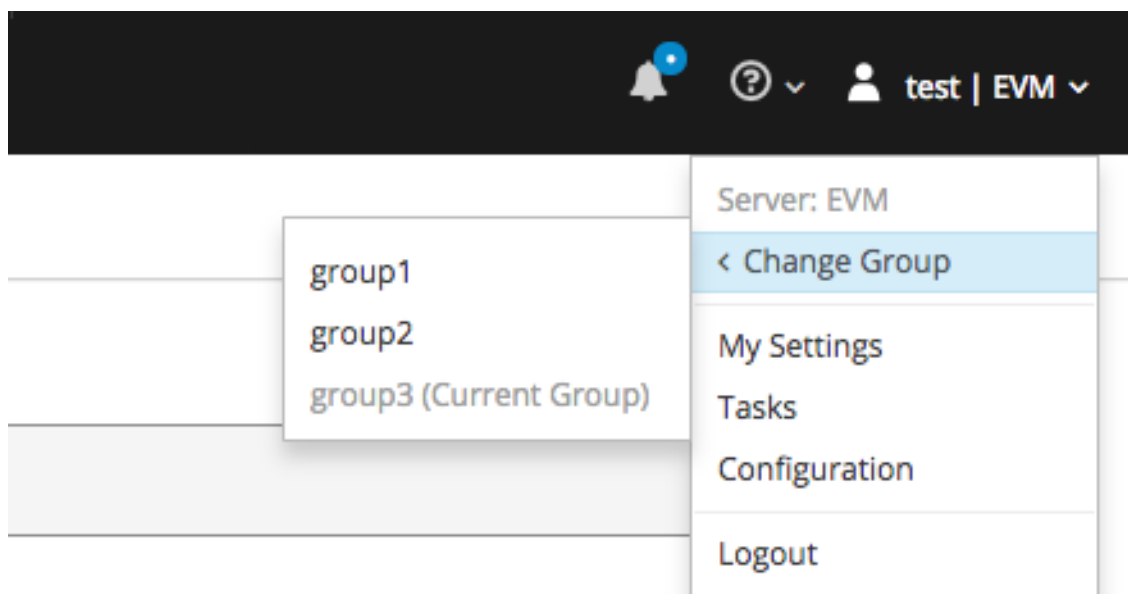
User groups create filters and assign roles to users. You can either create your own groups, or leverage your LDAP directory service to assign groups of users to account roles. For a list of what each pre-defined account role can do, see [Section 4.2.9, “Roles”](#).



NOTE



To assign account roles using LDAP groups, see [Section 4.1.4.2.4, “Assigning CloudForms Account Roles Using LDAP Groups”](#).

If a user belongs to multiple groups, you can change the user’s current group by navigating to the settings menu, clicking **Change Group**, then selecting the desired group. The **Current Group** affects quota and other settings for the user.



4.2.8. Creating a Group

To create a user group:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Groups**.
3. Click  (**Configuration**), and  (**Add a new Group**) to create a group.

4. Enter a name for the group in the **Name** field. To ensure compatibility with tags, use underscores in place of spaces. For example, Red Hat CloudForms-**test_group**.
5. Select a **Role** to map to this group.
6. Select the **Project/Tenant** this group must belong to.
7. Select any filters that you want applied to what this group can view in the **Assign Filters** area.
8. Check the boxes for the filters you want applied to this user. The items that have changed will show in blue italicized font.
9. Click the **Host & Clusters** tab.
10. Check the boxes for the host and clusters that you want to limit this user to. The items that have changed will show in blue italicized font.

Assign Filters

My Company Tags	Hosts & Clusters	VMs & Templates
-----------------	-----------------------------	-----------------

This user is limited to the selected items and their children.

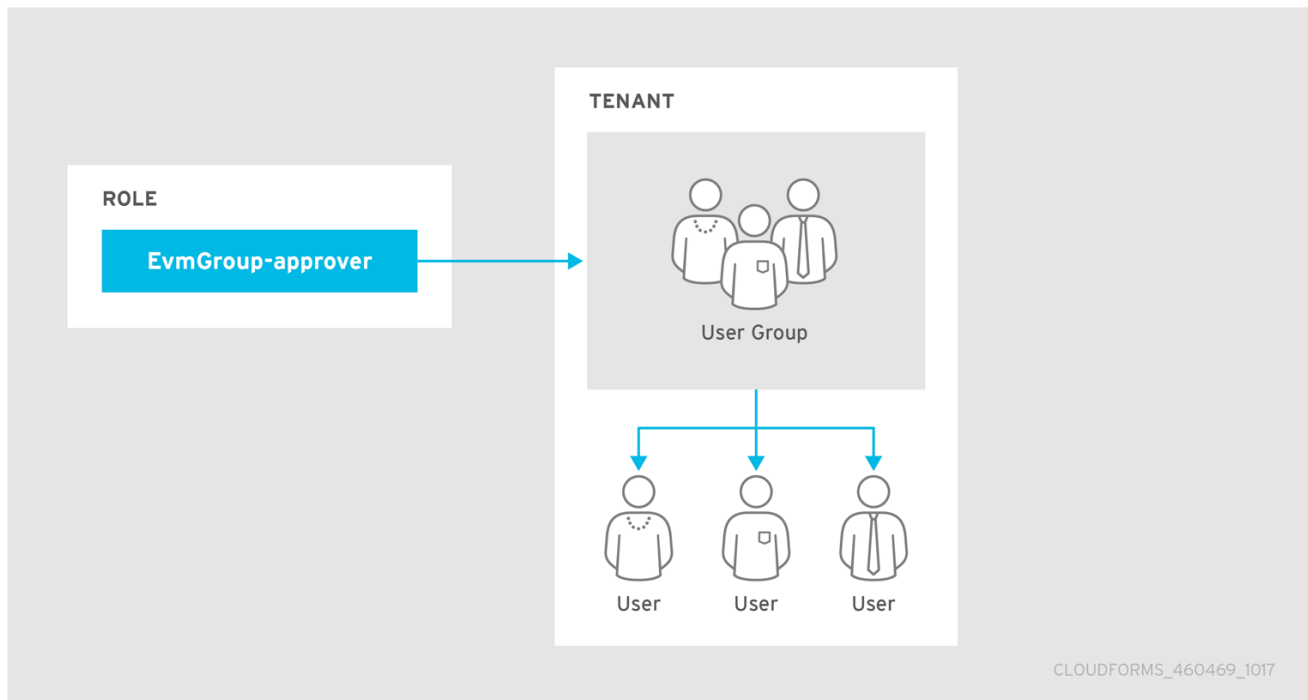


11. Click the **VMs & Templates** tab. This shows folders that you have created in your virtual infrastructure.
12. Check the boxes for the folders that you want to limit this user to. The items that have changed will show in blue italicized font.
13. Click **Add**.

After creating a group, assign a user to the group by editing a user.

4.2.9. Roles

When you create a group, you must specify a role to give the group rights to resources in the console. The group's role determines the scope of access for the users that are members of the group.



A group can only be assigned one role when the local database is used for authentication, as shown in the above image.

Red Hat CloudForms provides a default group of roles, but you can also create your own, or copy and edit the default groups.



NOTE

If you have enabled **Get Role from LDAP** under **LDAP Settings**, then the role is determined by the LDAP user's group membership in the directory service. See [Section 4.1.4.2.2, "Configuring LDAP Authentication"](#).

To view details of a role and its level of access:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Roles**.
3. Click on a role from the list to display role information and the product features the role can access (marked by a checkmark). You can expand the categories under **Product Features** to see further detail.

The table below shows a summary of the functions available to each role.

4.2.9.1. Account Roles and Descriptions



Role	Description
Administrator	Administrator of the virtual infrastructure. Can access all infrastructure functionality. Cannot change server configuration.

Role	Description
Approver	Approver of processes, but not operations. Can view items in the virtual infrastructure, view all aspects of policies and assign policies to policy profiles. Cannot perform actions on infrastructure items.
Auditor	Able to see virtual infrastructure for auditing purposes. Can view all infrastructure items. Cannot perform actions on them.
Container Administrator	Administrator with capabilities to configure, view and execute tasks on all containers and related underlying infrastructure. Has access to Nodes, Pods and Projects dashboards.
Container Operator	This role can view and execute tasks related to containers and related underlying infrastructure. The Container Operator has access to locked versions of the same dashboards as the Container Administrator.
Desktop	Access to VDI pages.
Operator	Performs operations of virtual infrastructure. Can view and perform all functions on virtual infrastructure items including starting and stopping virtual machines. Cannot assign policy, but can view policy simulation from Virtual Machine page.
Security	Enforces security for the virtual environment. Can assign policies to policy profiles, control user accounts, and view all parts of virtual infrastructure. Cannot create policies or perform actions on virtual infrastructure.
Super Administrator	Administrator of Red Hat CloudForms and the virtual infrastructure. Can access all functionality and configuration areas.
Support	Access to features required by a support department such as diagnostics (logs). Can view all infrastructure items and logs. Cannot perform actions on them.
Tenant Administrator	Configures settings applicable to a Tenant. Sets Branding, maps groups/roles, configures LDAP credentials, and configures dashboard settings.
Tenant Quota Administrator	Configures quota limits for the tenant, applying usage constraints for CPU, Memory, Storage, Maximum number of VMs, and Maximum number of Templates.
User	User of the virtual infrastructure. Can view all virtual infrastructure items. Cannot perform actions on them.

Role	Description
User Limited Self Service	Limited User of virtual machines. Can make provision requests. Can access some functions on the virtual machine that the user owns including changing power state.
User Self Service	User of virtual machines. Can make provision requests. Can access some functions on the virtual machine that the user owns and that the user's LDAP groups own including changing power state.
Vm User	User of virtual machines. Can access all functions on the virtual machine including changing power state and viewing its console. Cannot assign policy, but can view policy simulation from virtual machine page.

4.2.10. Creating a Role

To create a role:

1. From the settings menu, select **Configuration**.
2. Click on the **Access Control** accordion, then click **Roles**.
3. Click  (Configuration), and  (Add a new Role). Alternatively, you can copy an existing role to a new role by clicking **Copy this to a new Role**.
4. In the **Role Information** area, type a name for the new role. For **VM & Template Access Restriction**, select if you want to limit users with this role to only see virtual machines owned by the user, owned by the user or its group, or all virtual machines.

Adding a new Role

Role Information

Name

Evm-Role-user

VM & Template Access Restriction

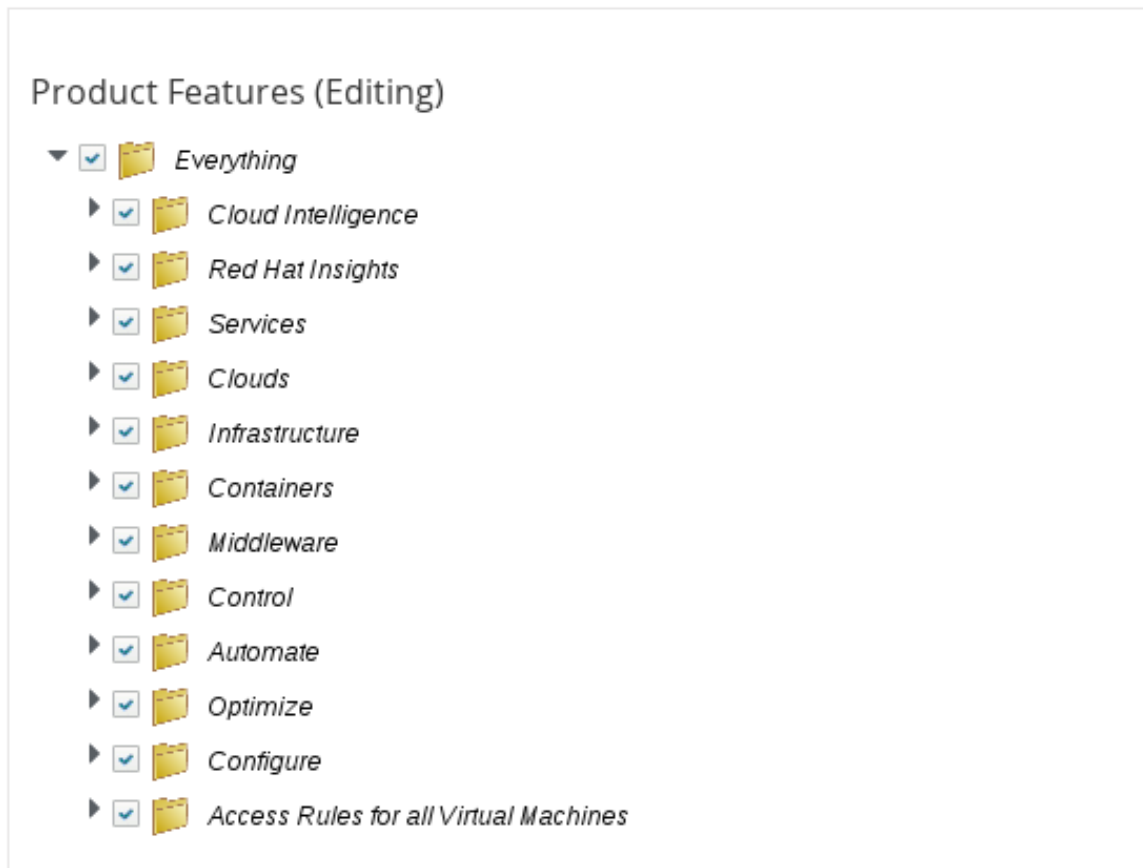
None

None

Only User or Group Owned

Only User Owned

- Under **Product Features (Editing)**, navigate to the appropriate feature and enable or disable it.



* Double click a feature to open/close all children.

- Click **Add**.

4.3. DIAGNOSTICS

From the settings menu, select **Configuration**. Click on the **Diagnostics** tab to see the status of the different Red Hat CloudForms roles and workers for each server, view and collect logs, and gather data if there are any gaps in capacity and utilization information. The Diagnostics area is designed in a hierarchy.

- At the **region** level, you can see replication status, backup the VMDB, and run garbage collection on the VMDB.
- At the **zone** level, you can see Red Hat CloudForms roles by servers and servers by roles. In addition, you can set log collection values for a specific zone, and collect gap data for capacity and utilization.
- At the **server** level, you can see the workers for each server, set log collection values for a specific server, and view current logs.

4.3.1. Region Diagnostics

Using the console, you can set the priority of server regional roles, review and reset replication, and create backups of your database either on demand or on a schedule.

Regions are used primarily to consolidate multiple VMDBs into one master VMDB for

reporting while zones are used to define functional groups of servers. There can be only one region per VMDB, but multiple zones per region (or VMDB). Some server roles are aware of each other across Red Hat CloudForms appliances at the region level. This means that redundancy and failover rules apply at the region level. You can also set priorities for the server roles that provide failover.

4.3.1.1. Server Role Priorities

If you have multiple servers in your environment with duplicate failover roles, then you can set the priority of the server role.

- Only server roles that support failover can be marked as primary. These roles only allow one server to be active at a time. These are **Notifier**, **Capacity & Utilization Coordinator**, **Event Monitor**, **Scheduler**, **Storage Inventory**, and **Provider Inventory**.
- All other server roles are additive. The more servers with that role in a zone the more work that can be performed.

There are three role priorities.

- **Primary:** There can only be one primary per zone or region per role. When an appliance is started, the system looks to see if any role is set to primary. If that is the case, the role is activated on that appliance and deactivated from the secondary. In the console, primary roles are shown in bold letters. The text turns red if the server goes down. You must actively set the primary priority.
- **Secondary:** This is the default priority. There can be multiple secondaries. When an appliance is started, if no primary is found in the zone, the first appliance to start takes the role. In the console, secondary roles are displayed normally with the word "secondary".
- **Tertiary:** If all appliances with primary roles or secondary roles were down, one of the tertiary would be activated. The reason for tertiary is to ensure that if a server with crucial roles such as Provider Inventory or Event Monitor goes down, you have a way to associate those roles to different appliances by organizing the priorities. Tertiary roles simply show as active in the console.





4.3.2. Region Aware Server Roles

Role	More than one per Region	Can have Priority Set
Automation Engine	Y	N
Database Operations	Y	N
Notifier	N	Y
Reporting	Y	N
Scheduler	N	Y

Role	More than one per Region	Can have Priority Set
User Interface	Y	N
Web Services	Y	N

4.3.3. Setting the Priority of a Failover Role

To set the priority of a failover role:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Depending on how you want to view your servers, click either the **Roles by Servers** tab or the **Servers by Roles** tab.
4. In the **Status** of **Roles for Servers** in **Zone Default Zone** area, click on the role that you want to set the priority for.
5. Click  (**Configuration**), and  (**Promote Server**) to make this the primary server for this role.
6. Click  (**Configuration**), and  (**Demote Server**) to demote the priority of this server for this role.

4.3.4. Zone Diagnostics

The console provides a way to see all the server roles that a server has been assigned and if these roles are running. This is especially helpful when you have multiple servers with different server roles. For each zone you can also set a central place for all logs to be collected, and collect capacity and utilization data that may be missing.

4.3.4.1. Viewing the Status of Server Roles

To view the status of server roles:


1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Depending on how you want to view your servers, click either **Roles by Servers** or the **Servers by Roles**.

4.3.4.2. Zone Aware Server Roles

Role	More than one per Region	Can have Priority Set
Automation Engine	Y	N
Capacity & Utilization Coordinator	N	Y
Capacity & Utilization Data Collector	Y	N
Capacity & Utilization Data Processor	Y	N
Database Operations	Y	N
Event Monitor	N	Y
Provider Inventory	N	Y
Provider Operations	Y	N
Notifier	N	Y
Reporting	Y	N
Scheduler	N	Y
SmartProxy	Y	N
SmartState Analysis	Y	N
User Interface	Y	N
Web Services	Y	N

4.3.4.2.1. Removing an Inactive Server

To remove an inactive server:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Click on the name of the server in the tree view.
4. Click  (**Delete Server**). This button is available only if the server is inactive.

4.3.4.3. Zone Log Collections

If you have multiple servers reporting to one central VMDB, then you can collect the configuration files and logs from the console of any of the servers. While you can set this

either at the zone or server level, settings at the server level supersede the ones at the zone level.



Log depot options include:

- Anonymous File Transfer Protocol (FTP)
- File Transfer Protocol (FTP)
- Network File System (NFS)
- Red Hat Dropbox
- Samba

See your network administrator if you need to set up one of these shares. You will also need a user that has write access to that location.

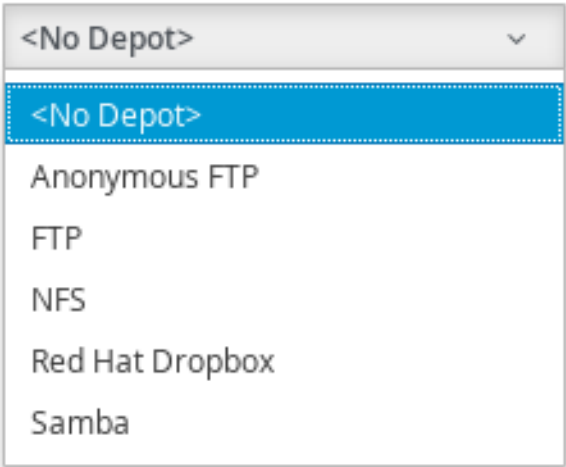
4.3.4.3.1. Setting the Location of the Log Depot

To set the location of the log depot:

1. From the settings menu, select **Configuration**.
2. Click the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Click **Collect Logs**.
4. Click  (**Edit**).
5. Select the **Type** of share. 
6. Using the fully qualified domain name (**FQDN**) of the depot server, type in the appropriate settings for the **URI**.

Editing Log Depot Settings for Zone: default

Type



A screenshot of a web application's dropdown menu for selecting a log depot type. The dropdown is open, showing a list of options. The top option, "<No Depot>", is highlighted with a blue background. Below it are the following options: "Anonymous FTP", "FTP", "NFS", "Red Hat Dropbox", and "Samba". The dropdown has a small downward arrow icon in the top right corner.

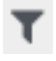
Type
<No Depot>
Anonymous FTP
FTP
NFS
Red Hat Dropbox
Samba

7. If required, enter your user **ID** and **password** then click **Validate** to confirm the settings.

8. Click **Save**.

4.3.4.3.2. Collecting and Downloading Logs from All Servers in a Zone

To collect and download logs from all servers in a zone:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Click the **Collect Logs** tab.
4. Click  (**Collect all logs**). All files in the logs directory as well as configuration files are collected from the selected zone.
5. Click **OK**. The status of the log retrieval shows in the Red Hat CloudForms console.

4.3.4.4. Capacity and Utilization Repair

Under certain circumstances, it is possible that Red Hat CloudForms is not able to collect capacity and utilization data. This could be due to password expiration, a change in rights to the cloud provider and this change didn't provide enough granularity to the Red Hat CloudForms service account, or network connectivity. The gap data is collected directly by extracting the monthly performance data. Gap collection need to be completed for each zone individually. Therefore, the procedure below need to be repeated for each zone.

4.3.4.4.1. Repairing Capacity and Utilization Data

To repair capacity and utilization data:

1. Log in to a Red Hat CloudForms appliance located in the zone for which you want to gather the data.
2. From the settings menu, select **Configuration**.
3. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
4. Click **C & U Gap Collection**.
 - a. Select the appropriate **Timezone**.



NOTE

Do not select more than one week unless instructed to do so by Red Hat Support.

- b. Select a **Start Date**.
 - c. Select an **End Date**.
5. Click **Submit**.

After the gap collection has completed for this zone, repeat these same steps for the next zone. You can check for completion by going to the clusters page and checking for the capacity and utilization data for the time period specified.

4.3.5. Server Diagnostics

Under **Diagnostics** for a server, you can view the status of Red Hat CloudForms workers running on the server, set log collection setting for only that server, and view the server's current Red Hat CloudForms and audit logs.

4.3.5.1. Workers

The **Workers** tab enables you to see the status of and restart Red Hat CloudForms workers.

You can see additional information on and restart the following items:

- **C & U Metrics Collectors** that collects capacity and utilization data.
- **C & U Metrics Processors**, which processes the collected capacity and utilization data.
- **Event Handlers** put events from the Event Monitor into the VMDB and starts Red Hat CloudForms processes if needed base on that information.
- **Event Monitors** that communicate with the external cloud provider to deliver up to date event information.
- **Generic Workers** that perform long running and priority processes.
- **Priority Workers** that perform high priority, short processes.
- **Schedule Workers** that maintains any items that run on a schedule.
- **Session Broker** that maintains a single connection to the cloud providers .
- **Refresh Workers** that runs the refresh processes.
- **Reporting Workers** that generate reports.
- **SmartProxy Workers** that run SmartState Analyses on virtual machine.
- **User Interface Worker** that allows users access to the console.
- **Web Services Worker** that maintains Red Hat CloudForms Web services.
- **VM Analysis Collectors** that run and process SmartState Analyses on virtual machines.

4.3.5.1.1. Reloading Worker Display



To reload worker display:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click the **Workers** tab.

5. Click  (**Refresh Current Workers display**).

4.3.5.1.2. Restarting a Worker

To restart a worker:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click on the **Workers** tab.
5. Click on the worker you want to restart.
6. Click  (**Configuration**), then  (**Restart selected worker**).
7. Click **OK**.

4.3.5.2. Server and Audit Logs

4.3.5.2.1. Collecting Server Logs and Configuration Files

While you can designate a central location to collect logs for all servers in a specific zone, you can override those values for a specific server. To do this, designate a log depot location to store the files.

Log depot options include:


- Anonymous File Transfer Protocol (FTP)
- File Transfer Protocol (FTP)
- Network File System (NFS)
- Red Hat Dropbox
- Samba

See your network administrator to set up one of these shares. You also need a user that has write access to that location. Settings at the server level supersede the ones at the zone level.

4.3.5.2.2. Setting the Location of the Log Depot for a Specific Server

To set the location of the log depot for a specific server

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to collect logs for.
4. Click on the **Collect Logs** tab.

5. Click  (**Edit Log Depot Settings for the selected Server**).
6. Select the **Type** of share.

Editing Log Depot Settings for Zone: default

Type



<No Depot> ▼

- <No Depot>
- Anonymous FTP
- FTP
- NFS
- Red Hat Dropbox
- Samba

7. Using the fully qualified domain name (**FQDN**) of the depot server, type in the appropriate settings for the **URI**.
8. Enter your user ID and password, then click **Validate** to confirm the settings.
9. Click **Save**.

4.3.5.2.3. Collecting the Current Log Set of a Server

To Collect the Current Log Set of a Server



1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to collect logs for.
4. Click on the **Collect Logs** tab.
5. Click  (**Collect**), then click  (**Collect current logs**). All current log files in as well as configuration files are collected.
6. Click **OK**. The status of the log retrieval shows in the Red Hat CloudForms console.

4.3.5.2.4. Collecting All Log Sets from a Server

To Collect All Log Sets from a Server

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to collect logs for.

4. Click **Collect Logs**.

5. Click  (**Collect**), then click  (**Collect all logs**). All files in the logs directory as well as configuration files are collected.

6. Click **OK**. The status of the log retrieval shows in the Red Hat CloudForms console.

4.3.5.2.5. Viewing the Server, Audit, and Production Logs

The server and audit logs roll over daily. The previous logs are stored as zipped files in the `/var/www/miq/vmdb/log` folder. The current logs can be easily viewed and downloaded from the settings menu; select **Configuration**, then click on the **Diagnostics** accordion.

Use the server log to see all actions taken by the server including communication with the SmartProxy and tasks.

4.3.5.2.6. Viewing the Server Log


To view the server log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **CFME Log**.

The Red Hat CloudForms server automatically retrieves the last 1000 lines of the log.

4.3.5.2.7. Reloading the Server Log

To reload the server log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **CFME Log**.
5. Click  (Reload the Log Display).

4.3.5.2.8. Downloading the Server Log

To download the server log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **CFME Log**.

5. Click  (**Download the Entire EVM Log File**).

**NOTE**

Use the **Audit Log** to see changes to the user interface and authentication.

4.3.5.2.9. Viewing the Audit Log


To view the audit log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **Audit Log**.

The server automatically retrieves the last 1000 lines of the log.


4.3.5.2.10. Reloading the Audit Log

To reload the audit log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **Audit Log**.
5. Click  (**Reload the Audit Log Display**).

4.3.5.2.11. Downloading the Audit Log

To download the audit log:

1. From the settings menu, select **Configuration**.
2. Click on the Diagnostics accordion, then click the Zone that you want to view.
3. Select the server that you want to view.
4. Click Audit Log.
5. Click  (**Download the Entire Audit Log File**).

4.3.5.2.12. Viewing the Production Log

Use the production log to see all actions performed using the console.


To view the production log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **Production Log**.

The Red Hat CloudForms server automatically retrieves the last 1000 lines of the log.


4.3.5.2.13. Reloading the Production Log

To reload the production log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **Production Log**.
5. Click  (**Reload the Product Log Display**).

4.3.5.2.14. Downloading the Production Log

To download the production log:

1. From the settings menu, select **Configuration**.
2. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Select the server that you want to view.
4. Click **Production Log**.
5. Click  (**Download the Production Log File**).

4.4. DATABASE OPERATIONS

4.4.1. Viewing Information on the VMDB

The **Database** accordion displays a summary of VMDB information, a list of all tables and indexes, settings for the tables, active client connection, and database utilization.

To view information on the VMDB:

1. From the settings menu, select **Configuration**.
2. Click the **Database** accordion.
3. Click **VMDB** in the tree view on the left.
4. Click the appropriate tab to view the desired information:

- **Summary:** displays statistics about the database.
- **Tables:** displays a clickable list of all the tables.
- **Indexes:** displays a clickable list of all the indexes.
- **Settings:** displays a list of all tables, their descriptions, and other valuable Information.
- **Client Connections:** displays all current connections to the VMDB.
- **Utilization:** displays usage charges for the disk and index nodes.

4.4.2. Database Regions and Replication

Regions are used to create a central database for reporting and charting. Do not use the database at the top level for operational tasks such as SmartState analysis or capacity and utilization data collection. Assign each server participating in the region a unique number during the regionalization process, then set your subordinate regions to replicate to the top region.



IMPORTANT

All Red Hat CloudForms databases in a multi-region deployment must use the same security key.

4.4.2.1. Creating a Region

In principle, a region is created when you set up your Red Hat CloudForms environment on the first appliance for the region. However, you can also create a region on an appliance where a database has already been set up. This process involves dropping and rebuilding the existing database to accommodate the new region number, and takes several minutes to complete.



WARNING

Performing this procedure destroys any existing data and cannot be undone. Back up the existing database before proceeding. By default, new appliances are assigned region 0. Do not use this number when creating a region because duplicate region numbers can compromise the integrity of the data in the database.

1. Log in to the appliance as the **root** user.
2. Enter **appliance_console**, and press **Enter**.
3. Press any key.
4. Select **Stop EVM Server Processes**.
5. Enter **Y** to confirm.

6. Choose **Configure Database**.
7. Enter a database region number that has not been used in your environment. Do not enter duplicate region numbers because this can corrupt the data.
8. Enter **Y** to confirm.
9. The menu reappears after all processes are complete.
10. Select **Start EVM Server Processes**.
11. Enter **Y** to confirm.
12. Select **Quit** to exit the advanced settings menu.

4.4.3. Configuring Database Replication and Centralized Administration

To configure database replication, you must configure one Red Hat CloudForms instance to act as a global copy, and one or more other instances to act as remote copies. This database replication relationship can only be configured on Red Hat CloudForms instances that are of the same version.



NOTE

Configuring database replication in this version of CloudForms automatically enables centralized administration, eliminating the need for further configuration.

Centralized administration in Red Hat CloudForms supports life cycle management operations that can be initiated from the global copy and processed and executed on the remote copy.



IMPORTANT

- You must configure at least one remote copy before you can configure the global copy. Changes to the role of a Red Hat CloudForms instance take several minutes to take effect.
- The region number must be unique on each Red Hat CloudForms instance where replication is configured. See [Section 4.4.2.1, “Creating a Region”](#) for instructions on how to create a region.

4.4.3.1. Configuring a Remote Copy

Configure a Red Hat CloudForms instance to act as a remote copy from which data will be replicated to the global copy.

1. Navigate to the settings menu.
2. Click **Configuration**.
3. Click the **Settings** accordion.
4. Click the region where the instance is located.

5. Click **Replication**.
6. Select **Remote** from the **Type** list.
7. Click **Save**.

4.4.3.2. Configuring the Global Copy

Configure a Red Hat CloudForms instance to act as the global copy to which data is replicated from the remote copies.

Details C & U Collection My Company Categories My Company Tags Import Tags Import Variables Map Tags Red Hat Updates **Replication**

Type **Global**

Subscriptions

At least 1 subscription must be added to save server replication type

Add Subscription

Database	Host	Username	Password	Port	Backlog	Actions
vmdb_production	10.64.15.100	root	5432		Accept ⋮

Save Reset

1. Navigate to the settings menu.
2. Click **Configuration**.
3. Click the **Settings** accordion.
4. Click the region where the instance is located.
5. Click **Replication**.
6. Select **Global** from the **Type** list.
7. Click **Add Subscription**.
 - a. Enter the name of the database on the remote copy in the **Database** field.
 - b. Enter the IP address or fully qualified domain name of the remote copy in the **Host** field.
 - c. Enter the name of a database user able to access the database in the **Username** field.
 - d. Enter the password of the database user in the **Password** field.
 - e. Enter the port by which the database is accessed in the **Port** field.
 - f. In **Actions**, click **Accept**. You can also **Update** your subscription if required, or click the menu on the right in order to **Validate** or **Discard**.
8. Click **Save**.

**NOTE**

Once you configure a Red Hat CloudForms instance to act as a global copy, and one or more other instances to act as remote copies, centralized administration is automatically enabled after the initial data sync is complete.

Database replication and centralized administration are now enabled on your instances. Centralized administration supports life cycle management operations, including virtual machine power operations and retirement, that can be initiated from the global copy and processed and executed on the remote copy.

See the section on *Centralized Administration* in the [Deployment Planning Guide](#) for an overview of the feature.

4.4.3.3. Resetting Database Replication

You can reset the replication relationship between the global copy and remote copies by temporarily removing and re-enabling the subscription from the global copy.

1. From the settings menu, select **Configuration**.
2. Click the **Settings** accordion.
3. Click the region where the instance is located.
4. Click **Replication**.
5. Remove the subscription:
 - a. Click the actions button for the subscription to reset.
 - b. Click **OK**.
6. Click **Save**.
7. Re-add the subscription:
 - a. Click **Add Subscription**.
 - b. Enter the name of the database on the remote copy in the **Database** field.
 - c. Enter the IP address or fully qualified domain name of the remote copy in the **Host** field.
 - d. Enter the name of a database user able to access the database in the **Username** field.
 - e. Enter the password of the database user in the **Password** field.
 - f. Enter the port by which the database is accessed in the **Port** field.
8. Click **Save**.

4.4.4. Backing Up and Restoring a Database

4.4.4.1. Running a Single Database Backup

To run a single database backup:

1. From the settings menu, select **Configuration**.
2. Click the **Diagnostics** accordion and click the **Region** name.
3. Click the **Database** tab.
4. If you have created a backup schedule and want to use the same depot settings, select the schedule in the **Backup Schedules** box.
5. If you do not want to use the settings from a backup schedule, select a type of server for storing the backups from the Type drop-down list in the **Database Backup Settings** box. You can use **Network File System** (NFS) or **Samba**.
 - If you select **Samba**, enter the **URI**, **User ID**, and a valid **Password**. Click **Validate** to check the settings.
 - If you select **Network File System**, enter the **URI**.
6. Click **Submit** to run the database backup.

4.4.4.2. Restoring a Database from a Backup

If a database is corrupt or fails, restore it from a backup. You can restore a backup from a local file, NFS, or Samba.

To restore a database from a backup:

1. Save the database backup file as **/tmp/evm_db.backup**. Red Hat CloudForms looks specifically for this file when restoring a database from a local backup.
2. If you are restoring a database backup on a high availability environment, stop the **regmgrd** service. This is not required in other CloudForms configurations.

```
# systemctl stop rh-postgresql95-repmgr
```
3. Log in to the appliance as the **root** user.
4. Enter **appliance_console**, and press **Enter**.
5. Select **Stop EVM Server Processes** to stop processes on all servers that connect to this VMDB.
6. Enter **Y** to confirm.
7. After all processes are stopped, press **Enter** to return to the menu.
8. Press **Enter** again to manually configure settings.
9. Select **Restore Database From Backup**, then specify the location to restore the backup from in the **Restore Database File** menu:
 - a. If you saved the database backup file locally as **/tmp/evm_db.backup**, select **Local file**. You can also restore from a **Network File System (nfs)** or **Samba (smb)**.

- b. Specify the location of the backup file.



NOTE

The appliance console menu may respond slowly if connections are open and the server is still shutting down. If this occurs, wait a minute and try again.

10. Enter **Y** to keep the database backup after restoring from it. Enter **N** to delete it.
11. Press **Y** to confirm.
12. After the backup completes, press **Enter** to return to the menu.
13. Press **Enter** again to manually configure settings.
14. Select **Start EVM Server Processes** to restart all processes on servers that connect to this VMDB.
15. Enter **Y** to confirm.
16. If you are restoring a database backup on a high availability environment, start the **regmgrd** service. This is not required in other CloudForms configurations.

```
# systemctl start rh-postgresql95-repmgr
```

4.4.5. Performing a Binary Backup and Restoring the Database

Preserve data at the file system level by performing a binary backup. This includes all databases, users and roles, and other objects.




NOTE

This procedure uses the **pg_basebackup** utility to perform a remote database backup and create a full replacement of the PostgreSQL data directory, capturing the exact state of the database when the backup finishes. For more information on the **pg_basebackup** utility, see the PostgreSQL documentation.

4.4.5.1. Performing a Binary Backup

Create a binary backup and store it as a **gzip** tar file inside the CloudForms backup directory.



IMPORTANT

PostgreSQL superuser or user with *Replication* permissions are required to perform this procedure.

1. SSH into the database server as the **root** user or provide PostgreSQL superuser credentials.
2. Run the **pg_basebackup** command to create the backup.

```
# pg_basebackup -x -h hostname -U root -Ft -z -D filename
```

■

where:

-h *hostname*

Specifies the IP address of the database server.

-D *filename*

Specifies the name of the directory created to contain the backup.

4.4.5.2. Restoring a Database from the Backup

Restore your PostgreSQL binary backup using the following steps. This process will require stopping both EVM and PostgreSQL services before restoring data.

1. Copy the existing backup to the target VM

```
# scp filename/base.tar.gz root@hostname:/var/www/miq
```

2. SSH to the target VM.

```
# ssh root@hostname
```

3. Stop both the EVM and PostgreSQL servers.

```
# systemctl stop evmserverd
# systemctl stop $APPLIANCE_PG_SERVICE
```

4. Rename the existing data directory.

```
# mv /var/opt/rh/rh-postgresql95/lib/pgsql/data /var/opt/rh/rh-
postgresql95/lib/pgsql/data.backup
```

5. Create a clean data directory.

```
# mkdir /var/opt/rh/rh-postgresql95/lib/pgsql/data
```

6. Unzip the tar file to the new directory.

```
# tar -xzf /var/www/miq/base.tar.gz -C /var/opt/rh/rh-
postgresql95/lib/pgsql/data
```

7. Correct permissions.

```
# chown postgres:postgres /var/opt/rh/rh-postgresql95/lib/pgsql/data
# chmod 700 /var/opt/rh/rh-postgresql95/lib/pgsql/data
```

8. Restart the PostgreSQL and EVM servers.

```
# systemctl start $APPLIANCE_PG_SERVICE
# systemctl start evmserverd
```

4.4.6. Running Database Garbage Collection

The database server collects garbage automatically, but Red Hat may occasionally direct you to run database garbage collection manually in order to reclaim unused space in your VMDB.

To run database garbage collection:

1. From the settings menu, select **Configuration**.
2. Click the **Diagnostics** accordion and click the **Region** name.
3. Click the **Database** tab.
4. In the **Run Database Garbage Collection Now** box, click **Submit**.

4.4.7. Changing the Database Password

Use the following procedures to change the password on appliances containing an internal database, and on worker appliances.

To change the password for an external database, use the procedure in [Section 4.4.7.2, “Changing the Password on the Worker Appliances”](#).



NOTE

See [Appliance Types](#) in the *Deployment Planning Guide* for a summary of different types of appliances.

4.4.7.1. Changing the Password on the Database Appliance

Red Hat CloudForms provides a default database password for the internal PostgreSQL database.

To change the password, you need to stop the CloudForms service, change the password for the PostgreSQL database, run a command to change the password in the configuration file that **evmserverd** uses to access the server, and restart the CloudForms appliance.

1. Stop the CloudForms service:
 - a. SSH into the appliance.
 - b. To stop the CloudForms service, run the following command:

```
service evmserverd stop
```

2. Use **pgadmin** to change the password for your CloudForms database (default is **vmdb_production**). If you do not have **pgadmin**, you can change the password by running:

```
psql -U root -d vmdb_production
```

- a. At the `vmdb#` prompt, type:

```
ALTER USER root WITH PASSWORD 'newpassword';
```

- b. To exit **psql**, type:

```
\q
```

3. Change the password in the configuration file that **evmserverd** uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseyaml --password  
newpassword
```

4. Restart the CloudForms service:

```
service evmserverd start
```

5. Verify that you can log in to the CloudForms console.

4.4.7.2. Changing the Password on the Worker Appliances

1. Stop the CloudForms service:

- a. SSH into the appliance.

- b. To stop the CloudForms service, run the following command:

```
service evmserverd stop
```

2. Change the password in the configuration file that **evmserverd** uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseyaml --password  
newpassword
```

3. Restart the CloudForms service:

```
service evmserverd start
```



IMPORTANT

In a high availability environment, if using the same PostgreSQL user for replication, you must also change the password in the `/var/lib/pgsql/.pgpass` file on every database node.

Additionally, if the password for the user being used for region-to-region replication is changing, users must also change the password in the replication Copy subscription screen. See [Section 4.4.3.2, “Configuring the Global Copy”](#).

4.4.8. Adding a New Appliance to an Existing Region with a Non-default Password

1. Create the new appliance.
2. Start the appliance, but do not go into any of the configuration options, instead **SSH** into the new appliance.
3. In the `/var/www/miq/vmdb` directory, create a file called **REGION**. Its only contents should be the number of the Region that it is joining. (You could also just copy the **REGION** file from the VMDB appliance.)
4. Edit the **database.yml** file in the `/var/www/miq/vmdb` directory. (You may want to save from the original.)
 - a. Replace the contents of the **"production"** section with the contents of the **"base"** section.
 - b. Edit the **"host"** parameter to match the IP of the appliance hosting the VMDB.
 - c. Save the new **database.yml**.
5. Run the following command to change the password in the configuration file that **evmserved** uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseyaml --password
newpassword
```

6. Restart the new worker appliance:

```
service evmserved restart
```

4.4.9. Configuring Scheduled Database Maintenance

You can schedule hourly or periodic database maintenance through the appliance console. Performing regular PostgreSQL database maintenance helps to maintain a more responsive Red Hat CloudForms environment.

Hourly database maintenance tasks, such as reindexing, are useful for highly active database tables such as metrics, workers, and servers.

You also may want to perform periodic database maintenance to truncate empty metrics tables and reorganize the database. Periodic maintenance can be configured to run hourly, daily, weekly, or monthly, at a specified hour and on a specified day.



NOTE

Periodic maintenance can impact appliance performance while it is running. Red Hat recommends scheduling periodic maintenance infrequently, and at off hours.

To configure hourly and periodic database maintenance:

1. Log in to the appliance as the **root** user.

2. Enter **appliance_console**, and press **Enter**.
3. Press any key.
4. Select **Configure Database Maintenance** to configure the automatic database maintenance schedule through a dialog.
 - a. For **Configure Hourly Database Maintenance?** Type **y** or **n**.
 - b. For **Configure Periodic Database Maintenance?** Type **y** or **n**.

The next options depend on the periodic database maintenance frequency you choose, and are specified using the same dialog. The dialog finishes configuration with a "**Database maintenance configuration updated**" message when complete.

To reset your database maintenance settings, enter **Configure Database Maintenance** again from the appliance console menu, and confirm that you want to unconfigure the settings in the configuration dialog. This deletes the current settings.

To configure a new database maintenance schedule, enter the **Configure Database Maintenance** menu item once again and configure the values using the dialog.

CHAPTER 5. SMARTPROXIES

The embedded SmartProxy can analyze virtual machines that are registered to a host and templates that are associated with a provider.

5.1. INSTALLING THE SMARTPROXY FROM THE CONSOLE

The server comes with one SmartProxy version already available. It can also be installed on an ESX Server version 3.0.2, 3.5 or 4.



IMPORTANT

Contact Red Hat before installing a new SmartProxy on an ESX Server.



Requirements:

- On ESX, SSH (Secure Shell) must be enabled. This is usually port 22.
- 300 MB free disk space to install and run the SmartProxy.
- Administrator or root credentials.
- The host must already be in the VMDB either by discovery or manually.

5.2. ENTERING CREDENTIALS AND OPERATING SYSTEM FOR THE TARGET HOST

Set the credentials and operating system for the target host to prepare for the installation of SmartProxy.

To Enter Credentials and Operating System for the Target Host:

1. Navigate to **Compute → Infrastructure → Hosts**.
2. Select the host you want to edit.
3. Click  (**Configuration**), then  (**Edit this item**).
4. In **Credentials**, click the **Default** tab and enter your login credentials. If you are using domain credentials, the format for User ID must be in the format of `<domainname>\<username>`. For ESX hosts, if SSH login is disabled for the default user, click the **Remote Login** tab and enter a user with remote login access.

Credentials

Default

Remote Login

Web Services

IPMI

Username

root


Password

••••••••

Confirm Password

••••••••

Validate

**IMPORTANT**


If the target is a Windows host, disconnect all network connections between the Windows proxy and the target. If an existing connection uses a different set of credentials than those set in the console, the installation may fail.

5. Click **Validate** to verify the credentials.
6. If you added the host manually instead of **Host Discovery** or **Provider Refresh** finding it, select the host's operating system from the **Host Platform** drop-down box to ensure the host platform is available.
7. Click **Save**.

When remotely installing on Windows hosts, the SmartProxy file is first copied to a Windows proxy. That computer then installs the file to the target host. The Windows proxy is the same as when you select the Default Repository SmartProxy box. You can locate this by navigating to the settings menu and selecting **Configuration**, then clicking on the desired server, then the **Server** tab, and exploring the **Server Control** area.

CHAPTER 6. ABOUT



Click the question mark icon () at the top right of the Red Hat CloudForms user interface to open a dropdown menu with links to Red Hat CloudForms documentation and the Red Hat Customer Portal, and for information about the current CloudForms session.

- Click **Documentation** to open a list of guides. Click the title of a document to view it as a PDF in a new web browser tab.
- Click **Red Hat Customer Portal** to open the Red Hat Customer Portal in a new web browser tab.
- Click **About** to open a pop-up window showing information about the CloudForms version and the session of the logged-in user.

CHAPTER 7. RED HAT INSIGHTS

Red Hat Insights is a service that uses the collective knowledge to help end users proactively diagnose systems and avoid critical downtime situations. Red Hat Insights does this by having systems periodically check in similar to Red Hat Subscription Management.

Red Hat Insights provides an easy to use dashboard that enables system administrators and IT operations managers to quickly identify key risks to stability, security, or performance. A glance at the display allows users to sort by category, view details of the impact and resolution, and then quickly determine what systems are affected.



NOTE

Red Hat Insights is available as a technology preview in this release of Red Hat CloudForms. For more information on the support scope for features marked as technology previews, see [Technology Preview Features Support Scope](#)

Advantages of Red Hat Insights:

- Proactively solve problems and avoid downtime related to security exploits, performance degradation, and stability failures
- Feel confident in resolutions founded on 20,000 pieces of knowledge from certified Red Hat engineers
- Reduces man hours spent researching how to identify and resolve issues as IT budgets shrink

The Red Hat Insights plugin has the following options:

- Overview
- Rules
- Systems

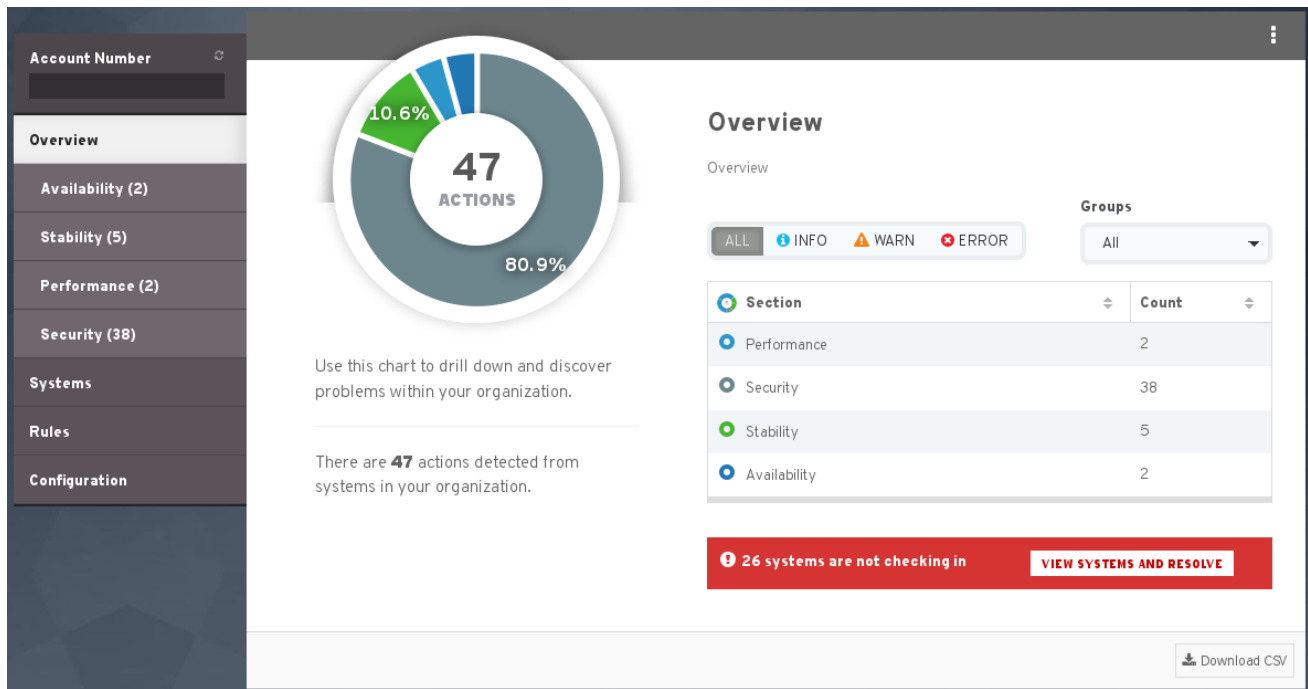
The following sections describe each of these tabs in more detail.

7.1. OVERVIEW TAB

The Overview tab helps you discover the issues within your deployment. It differentiates the issues under the following categories:

- Performance
- Security
- Stability
- Availability

You can select each of these categories to view more details on the issues.



7.2. RULES TAB

Rules enable easy addition of rules which operate on customer uploaded archives. It allows developers to focus on a single archive at a time while being able to process large amounts of data.

7.2.1. States

For management purposes, rules may be placed under one of four active states:

- **Active** - Rules which have been pushed to the master branch, are in prod and the content has been approved. This is the only state where the rules are displayed to customers.
- **Needs Content** - Plugins that the system has identified are in our master branch, hits have been found but do not have an entry / content written for them.
- **Inactive** - Once a rule is created from Needs Content, it will by default move to Inactive. Inactive can be used as a staging area as rules are written or to temporarily remove an active rule from the customer's view if further work needs to be completed. Rules can be deleted in this state
- **Retired** - Plugins or error info entries which are no longer in use. - Rules can be deleted in this state



7.2.2. Info Listed

You will find the following information available for each rule on the list

- **Error Key** - The returned key the plugin provides to alert detection.
- **Plugin** - Name of the plugin located in the master branch, ex: plugin.swappiness == plugins/swappiness.py

- **Description** - The 50 character "title" the customer sees during the drill down of issues.
- **Category** - Security, Stability, Performance.
- **Severity** - Warn, Error, Info.
- **Count** - The current amount of hosts the plugin has been detected by when the rules last ran.

7.3. SYSTEMS TAB

The Systems tab helps you discover the issues within your system. This tab lists the hostname of the system, the time of last check in and the status. You can filter the list by using  (**Actions**) for all systems that require actions and  (**No Actions**) for all systems that are working without issues and require no actions. You can also filter the list using the **Groups** dropdown list.

Account Number

Overview

Systems

Rules

Configuration

Systems

24 Systems with actions

37 Systems with no actions

Filter by System Actions

ALL SYSTEMS

WITH ACTIONS

WITHOUT ACTIONS

Groups

All

Show only systems not checking-in

Hostname	Last Check In	Status
Filter		
<input type="checkbox"/> dhcp210-2.gsslab.pnq.redhat.com	about an hour ago	
<input type="checkbox"/> dhcp233-127.gsslab.pnq.redhat.com	about 2 hours ago	
<input type="checkbox"/> rhel7-h.hkg.redhat.com	about 4 hours ago	
<input type="checkbox"/> rhev-m.hkg.redhat.com	about 4 hours ago	
<input type="checkbox"/> unused	about 4 hours ago	
<input type="checkbox"/> rhel-h2.hkg.redhat.com	about 4 hours ago	
<input type="checkbox"/> rhel7-satellite6.hkg.redhat.com	about 4 hours ago	
<input type="checkbox"/> mhuth-laptop	about 6 hours ago	