



Red Hat CloudForms 4.0

General Configuration

A guide to configuring and tuning CloudForms Management Engine

Red Hat CloudForms 4.0 General Configuration

A guide to configuring and tuning CloudForms Management Engine

Red Hat CloudForms Documentation Team
cloudforms-docs@redhat.com

Legal Notice

Copyright © 2017 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, registration, and updates. Information and procedures in this book are relevant to CloudForms Management Engine administrators. 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	8
2.2. DEFAULT VIEWS	9
2.2.1. Setting Default View for Management Engine	9
2.2.2. Setting Default Views for Infrastructure Components	10
2.2.3. Setting Default Views for Clouds	11
2.2.4. Setting Default Views for Containers	11
2.2.5. Setting Default Views for Services	13
2.3. DEFAULT FILTERS	13
2.3.1. Setting Default Filters for Cloud	13
2.3.2. Setting Default Filters for Containers	13
2.3.3. Setting Default Filters for Infrastructure	14
2.3.4. Setting Default Filters for Services	14
2.4. TIME PROFILES	14
2.4.1. Creating a Time Profile	14
2.4.2. Editing a Time Profile	15
2.4.3. Copying a Time Profile	16
2.4.4. Deleting a Time Profile	16
CHAPTER 3. CONFIGURATION	17
3.1. SETTINGS	17
3.1.1. Regions	17
3.1.1.1. Region Scope	20
3.1.1.2. Region Settings	20
3.1.1.3. Capacity and Utilization Collections	20
3.1.1.3.1. Capacity and Utilization Collection Settings	20
3.1.1.3.2. Enabling a Cluster, Host, or Datastore for Capacity and Utilization Collection	21
3.1.1.4. Tags	21
3.1.1.4.1. Company Tag Categories and Tags	21
3.1.1.4.2. Creating a Tag Category	21
3.1.1.4.3. Deleting a Tag Category	23
3.1.1.4.4. Creating a Company Tag	23
3.1.1.4.5. Deleting a Company Tag	23
3.1.1.4.6. Importing Tags for Virtual Machines	24
3.1.1.4.7. Importing Tags for a Virtual Machine from a CSV File	24
3.1.1.4.8. Importing Custom Values for Virtual Machines and Hosts	24
3.1.1.4.9. Importing Asset Tags for a Virtual Machine from a CSV File	25
3.1.1.5. Registering and Updating Red Hat CloudForms	26
3.1.1.5.1. Registering Appliances	26
3.1.1.5.2. Updating Appliances	28
3.1.1.5.3. Updating the Appliance Operating System	29
3.1.1.5.4. Subscription Management for Virtual Environments	30
3.1.2. Profiles	30
3.1.2.1. Creating an Analysis Profile	30
3.1.2.2. Creating a Host Analysis Profile	31
3.1.2.3. Creating a Virtual Machine Analysis Profile	32

3.1.2.4. Editing an Analysis Profile	33
3.1.2.5. Copying an Analysis Profile	33
3.1.2.6. Setting a Default Analysis Profile	33
3.1.3. Zones	34
3.1.3.1. Creating a Zone	34
3.1.3.2. Deleting a Zone	35
3.1.3.3. Editing a Zone	36
3.1.4. Servers	36
3.1.4.1. Changing Server Settings	36
3.1.4.1.1. Basic Information Settings	36
3.1.4.1.2. Server Control Settings	37
3.1.4.1.3. Server Roles	38
3.1.4.1.4. VMware Console Settings	40
3.1.4.1.5. NTP Servers Settings	41
3.1.4.1.6. Configuring SNMP	41
3.1.4.1.7. Outgoing SMTP Email Settings	42
3.1.4.1.8. Web Services Settings	43
3.1.4.1.9. Logging Settings	44
3.1.4.1.10. Custom Support URL Settings	44
3.1.4.2. Authentication	44
3.1.4.2.1. Changing an Authentication Setting	44
3.1.4.2.2. Authentication Settings	45
3.1.4.2.3. LDAP Settings	45
3.1.4.2.4. LDAP Role Settings	46
3.1.4.2.5. Amazon Settings	47
3.1.4.2.6. Trusted Forests	47
3.1.4.2.7. Adding Settings for a Trusted Forest	47
3.1.4.2.8. External Authentication (httpd)	48
3.1.4.3. Workers	51
3.1.4.3.1. Changing Settings for a Worker	51
3.1.4.3.2. Changing Settings for the Replication Worker	51
3.1.4.4. Database	52
3.1.4.4.1. Changing a Database Setting	52
3.1.4.5. Customization and Logos	53
3.1.4.5.1. Custom Logos	53
3.1.4.5.2. Uploading a Custom Logo to the Console	53
3.1.4.5.3. Customizing the Login Background	54
3.1.4.5.4. Customizing the Login Panel Text	55
3.1.4.6. Advanced Settings	55
3.1.4.6.1. Editing Configuration Files Manually	55
3.1.4.6.2. Configuration Parameters	56
3.1.5. Schedules	76
3.1.5.1. Scheduling SmartState Analyses and Backups	77
3.1.5.1.1. Scheduling a SmartState Analysis or Compliance Check	77
3.1.5.2. Scheduling a Database Backup	79
3.1.5.2.1. Modifying a Schedule	80
3.2. ACCESS CONTROL	80
3.2.1. Creating a Tenant	81
3.2.2. Creating a Project	81
3.2.3. Creating a Project Quota	81
3.2.4. Tagging Tenants and Projects	82
3.2.5. Creating a User	82
3.2.6. Deleting a User	83

3.2.7. Groups	83
3.2.8. Creating a User Group	83
3.2.9. LDAP Groups	84
3.2.10. Using CloudForms Management Engines Named Groups to Assign Account Roles	85
3.2.11. Account Role and Directory Service Group Names	85
3.2.12. Using Pre-existing LDAP Groups to Assign Account Roles	86
3.2.13. Roles	87
3.2.13.1. Account Roles and Descriptions	87
3.2.14. Creating a Role	88
3.3. DIAGNOSTICS	89
3.3.1. Region Diagnostics	90
3.3.2. Region Aware Server Roles	90
3.3.3. Setting the Priority of a Failover Role	91
3.3.4. Replication	91
3.3.4.1. Monitoring Database Replication	91
3.3.4.2. Resetting Replication	92
3.3.4.3. Running a Single Backup	92
3.3.4.4. Restoring a Database from Backup	93
3.3.5. Zone Diagnostics	93
3.3.5.1. Viewing the Status of Server Roles	93
3.3.5.2. Setting Server Role Priorities	93
3.3.5.3. Zone Aware Server Roles	94
3.3.5.4. Setting the Priority of a Failover Role	95
3.3.5.4.1. Removing an Inactive Server	95
3.3.5.5. Zone Log Collections	95
3.3.5.5.1. Setting the Location of the Log Depot	96
3.3.5.5.2. Collecting and Downloading Logs from All Servers in a Zone	96
3.3.5.6. Capacity and Utilization Repair	97
3.3.5.6.1. Repairing Capacity and Utilization Data	97
3.3.6. Server Diagnostics	98
3.3.6.1. Workers	98
3.3.6.1.1. Reloading Worker Display	98
3.3.6.1.2. Restarting a CloudForms Management Engine Worker	99
3.3.6.2. Server and Audit Logs	99
3.3.6.2.1. Collecting Server Logs and Configuration Files	99
3.3.6.2.2. Setting the Location of the Log Depot for a Specific Server	99
3.3.6.2.3. Collecting the Current Log Set of a Server	100
3.3.6.2.4. Collecting All Log Sets from a Server	101
3.3.6.2.5. Viewing the Server, Audit, and Production Logs	101
3.3.6.2.6. Viewing the Server Log	101
3.3.6.2.7. Reloading the Server Log	101
3.3.6.2.8. Downloading the Server Log	102
3.3.6.2.9. Viewing the Audit Log	102
3.3.6.2.10. Reloading the Audit Log	102
3.3.6.2.11. Downloading the Audit Log	103
3.3.6.2.12. Viewing the Production Log	103
3.3.6.2.13. Reloading the Production Log	103
3.3.6.2.14. Downloading the Production Log	103
3.4. DATABASE OPERATIONS	104
3.4.1. Viewing Information on the VMDB	104
3.4.2. Database Regions and Replication	104
3.4.2.1. Database Regions and Replication	104
3.4.2.2. Creating a Region	105

3.4.3. Replicating a Database	106
3.4.4. Enabling the Database Synchronization Role	106
3.4.4.1. Scheduling a Database Backup	107
3.4.4.2. Running a Single Database Backup	109
3.4.4.3. Restoring a Database from a Backup	109
3.4.4.4. Running Database Garbage Collection	110
3.4.5. Changing Database Password	110
3.4.5.1. Changing the Password on the Database Appliance	110
3.4.5.2. Changing the Password on the Worker Appliances	111
3.4.5.3. Adding a New Appliance to an Existing Region with Non-default Password	112
CHAPTER 4. SMARTPROXIES	113
4.1. INSTALLING THE SMARTPROXY FROM THE CONSOLE	113
4.2. ENTERING CREDENTIALS AND OPERATING SYSTEM FOR THE TARGET HOST	113
CHAPTER 5. ABOUT	115
CHAPTER 6. RED HAT ACCESS INSIGHTS	116
6.1. OVERVIEW TAB	116
6.2. RULES TAB	117
6.2.1. States	117
6.2.2. Info Listed	117
6.3. SYSTEMS TAB	118

CHAPTER 1. SETTINGS OVERVIEW

To view or modify global settings for your appliance, hover over the **Configure** menu and click the sub-tab **My Settings** to modify. The availability of each sub-tab depends on the role assigned to your user account. For more information on roles, see section **Account Roles and Descriptions**.

The following is a list of the sub-tabs available from the **Configure** menu:

- **My Settings** is available to all CloudForms Management Engine users. Its settings control the visual aspects of the console, time profiles, and tags used by the individual user.
- **Tasks** is used 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** is used to specify enterprise, region, zone, and server settings for the CloudForms Management Engine infrastructure. Diagnostics including logs and process status is also shown here.
- **About** provides session information and links to CloudForms Management Engine documentation as well as the Red Hat Customer Portal.

CHAPTER 2. MY SETTINGS

Options under **Configure** → **My Settings** enable you to control user settings such as how things are displayed, default views, and individual tags. You can also set your color scheme, button options, and external RSS feeds on the main **CloudForms Management Engine** dashboard.

2.1. VISUAL SETTINGS

For all of the **Visual** options, click **Save** to update your configuration settings. 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.

Grid/Tile Icons

Show Infrastructure Provider Quadrants	<input checked="" type="checkbox"/>
Show Cloud Provider Quadrants	<input checked="" type="checkbox"/>
Show Host Quadrants	<input checked="" type="checkbox"/>
Show Datastore Quadrants	<input checked="" type="checkbox"/>
Show VM Quadrants	<input checked="" type="checkbox"/>
Show Template Quadrants	<input checked="" type="checkbox"/>
Truncate Long Text	Middle (AB...34) ▴ ▾

- Check **Show Infrastructure Quadrants** to see the four icons in your provider. Uncheck to see only one icon.
- Check **Show Cloud Provider Quadrants** to see the four icons in your hosts. Uncheck to see only one icon.
- Check **Show Host Quadrants** to see the four icons in your hosts. Uncheck to see only one icon.
- Check **Show Datastore Quadrants** to see the four icons in your Datastores. Uncheck to see only one icon.
- Check **Show VM Quadrants** to see the four icons in your virtual machines. Uncheck to see only one icon.
- Check **Show Template Quadrants** to see the four icons in your templates. Uncheck to see only one icon.
- 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.

Use the following procedure to change grid and tile icons:


1. Navigate to **Configure** → **My Settings**, then click on the **Visual** tab.

2. In **Grid/Tile Icons**, select the items to display all four quadrants for.
3. Click **Save**.

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. Navigate to **Configure** → **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 **CloudForms Management Engine** dashboard, you can set the default start page to see a list of your virtual machines.

1. Navigate to **Configure** → **My Settings**, then click on the **Visual** tab.
2. In the **Start Page** area, select the page you want to see 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. Navigate to **Configure** → **My Settings**, then click on the **Visual** tab.
2. Make selections from **Display Settings** for the following items.

3. Use **Chart Theme** to select a group of colors and font sizes specifically for charts.
4. Use **Time Zone**



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).] to select the time zone in which to display the console.

5. Use **Locale** to select the language in which to display the console.
6. Click **Save**.

2.2. DEFAULT VIEWS







You can decide on the default views for your virtual machines, infrastructure, and other pages where the view is customizable. These settings can also be controlled on the actual pages where the items appear.


2.2.1. Setting Default View for Management Engine

Use the following procedure to set general view options:

1. Navigate to **Configure** → **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.



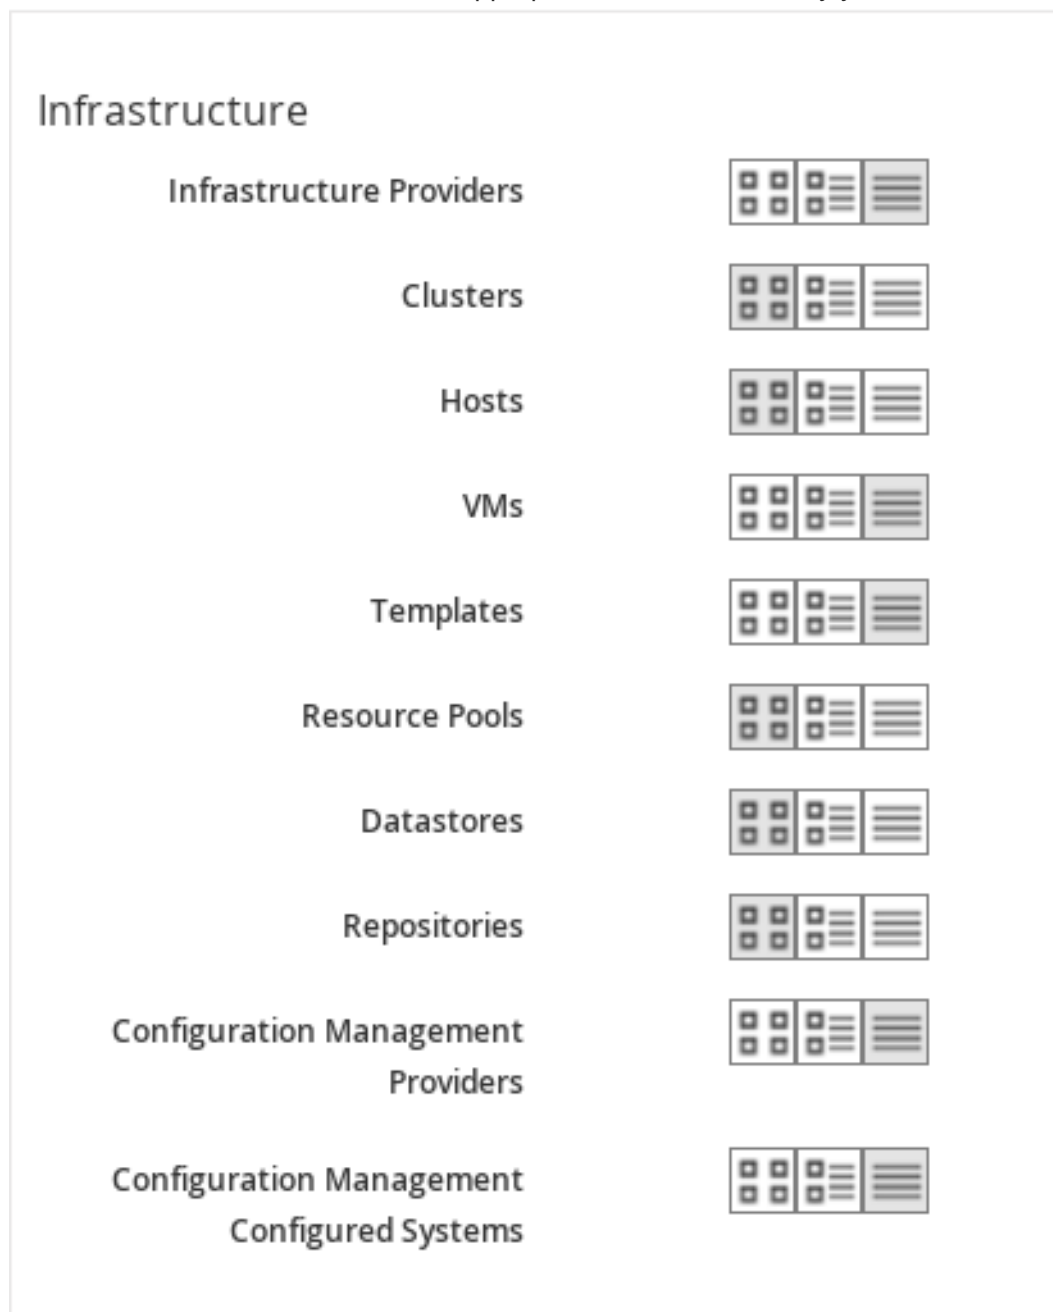
3. Click  (**Grid View**) to view virtual thumbnails or icons.
4. Click  (**Tile View**) for a view that combines the virtual thumbnail with some text properties that describe the items.
5. Click  (**List View**) or  (**Details Mode**) to view a detailed text listing.
6. Click  (**Expanded View**) for 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 Infrastructure Components

Use the following procedure to set default views for Infrastructure Components.

1. Navigate to **Configure** → **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**.

2.2.3. Setting Default Views for Clouds

Use the following procedure to set default views for clouds.

1. Navigate to **Configure** → **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.

Clouds

Cloud Providers



Availability Zones



Tenants



Flavors



Instances






Images



Stacks



- 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 Containers





Use the following procedure to set default views for services.

1. Navigate to **Configure** → **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.

Containers

Containers Providers	  
Nodes	  
Pods	  
Services	  
Routes	  
Projects	  
Replicators	  
Images	  
Image Registries	  
Containers	  
Container Topologies	  
Container Dashboards	  

- 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  (**Detail View**) to view a detailed text listing.
- Click  (**List View**) to view a text listing.





3. Click **Save**.

2.2.5. Setting Default Views for Services

Use the following procedure to set default views for services.

1. Navigate to **Configure** → **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  (Detail View) to view a detailed text listing.
- Click  (List View) to view a text listing.

3. Click **Save**.

2.3. DEFAULT FILTERS

You can set 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:

1. Navigate to **Configure** → **My Settings**, then click on the **Default Filters** tab.
2. In the **Hosts** folder, select the default filters that you want available on the **Hosts** page. Items that have changed show in blue, bold text.
3. From the **Cloud** folder, check the boxes for the default filters that you want available. Items that have changed show in blue and bold text.
4. Click **Save**.

2.3.2. Setting Default Filters for Containers

To Set Default Filters for containers:

1. Navigate to **Configure** → **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 and bold text.
3. Click **Save**.

2.3.3. Setting Default Filters for Infrastructure

To Set Default Filters for Infrastructure:

1. Navigate to **Configure** → **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, bold text. Not all filters are listed in the figure below.
3. Click **Save**.

2.3.4. Setting Default Filters for Services

To Set Default Filters for Services:

1. Navigate to **Configure** → **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, bold text. Not all filters are listed in the figure below.
3. Click **Save**.

2.4. TIME PROFILES

Time profiles limit the hours for which data is displayed when viewing capacity and utilization screens. They are also used for performance and trend reports, and for **Optimize** pages.

2.4.1. Creating a Time Profile

To Create a Time Profile:

1. Navigate to **Configure** → **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 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) American Samoa ▼													
Roll Up Daily Performance	<input type="checkbox"/>													

Add Cancel

- Type a meaningful name in the **Description** field.
- For **Scope**, select **All Users** to create a global time profile 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.
- Check the **Days** and **Hours** for the time profile.
- For **Timezone**, you can select a specific time zone or, you can let the user select a time zone when displaying data.
- 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.
- 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. Navigate to **Configure** → **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. Navigate to **Configure** → **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. Navigate to **Configure** → **My Settings**, then click on the **Time Profiles** tab.
2. Check the time profile you want to edit.
3. Click  (**Configuration**), and  (**Delete Selected Time Profiles**).
4. Make the required changes.
5. Click **Save**.

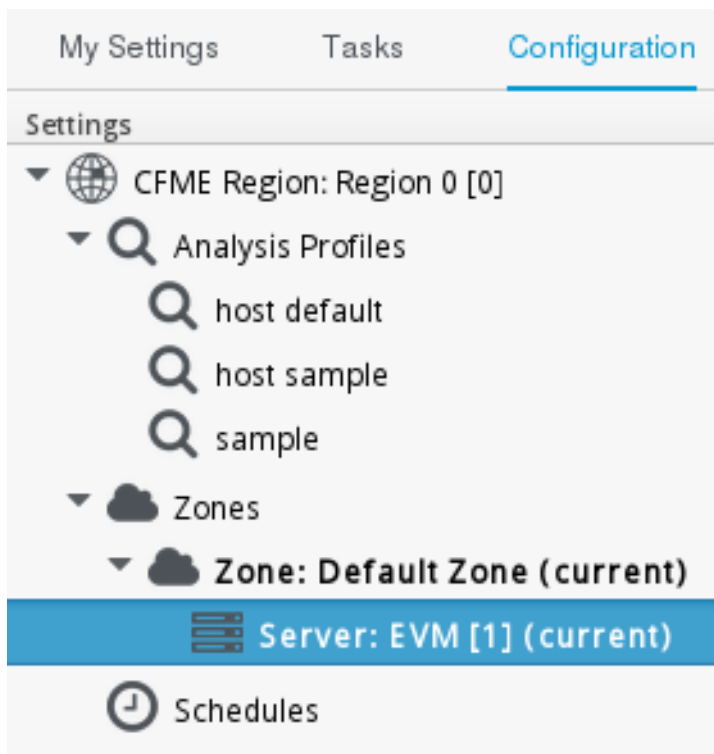
CHAPTER 3. CONFIGURATION

From the **Configuration** area, you can specify operating parameters for the CloudForms Management Engine infrastructure, view diagnostic information, and analytics on the servers. The accordion menu shows your CloudForms Management Engine infrastructure at the enterprise, zone, and server levels. There are three main areas.

- **Settings** enable you to modify the configuration of your CloudForms Management Engine infrastructure. You can also create analysis profiles and schedules for these profiles.
- **Diagnostics** displays the status of your servers and their roles and provides access to logs.

3.1. SETTINGS

Under **Configure** → **Configuration**, then in the **Settings** accordion, you have a hierarchy of the configurable items in your CloudForms Management Engine architecture. 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**.



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

3.1.1. Regions

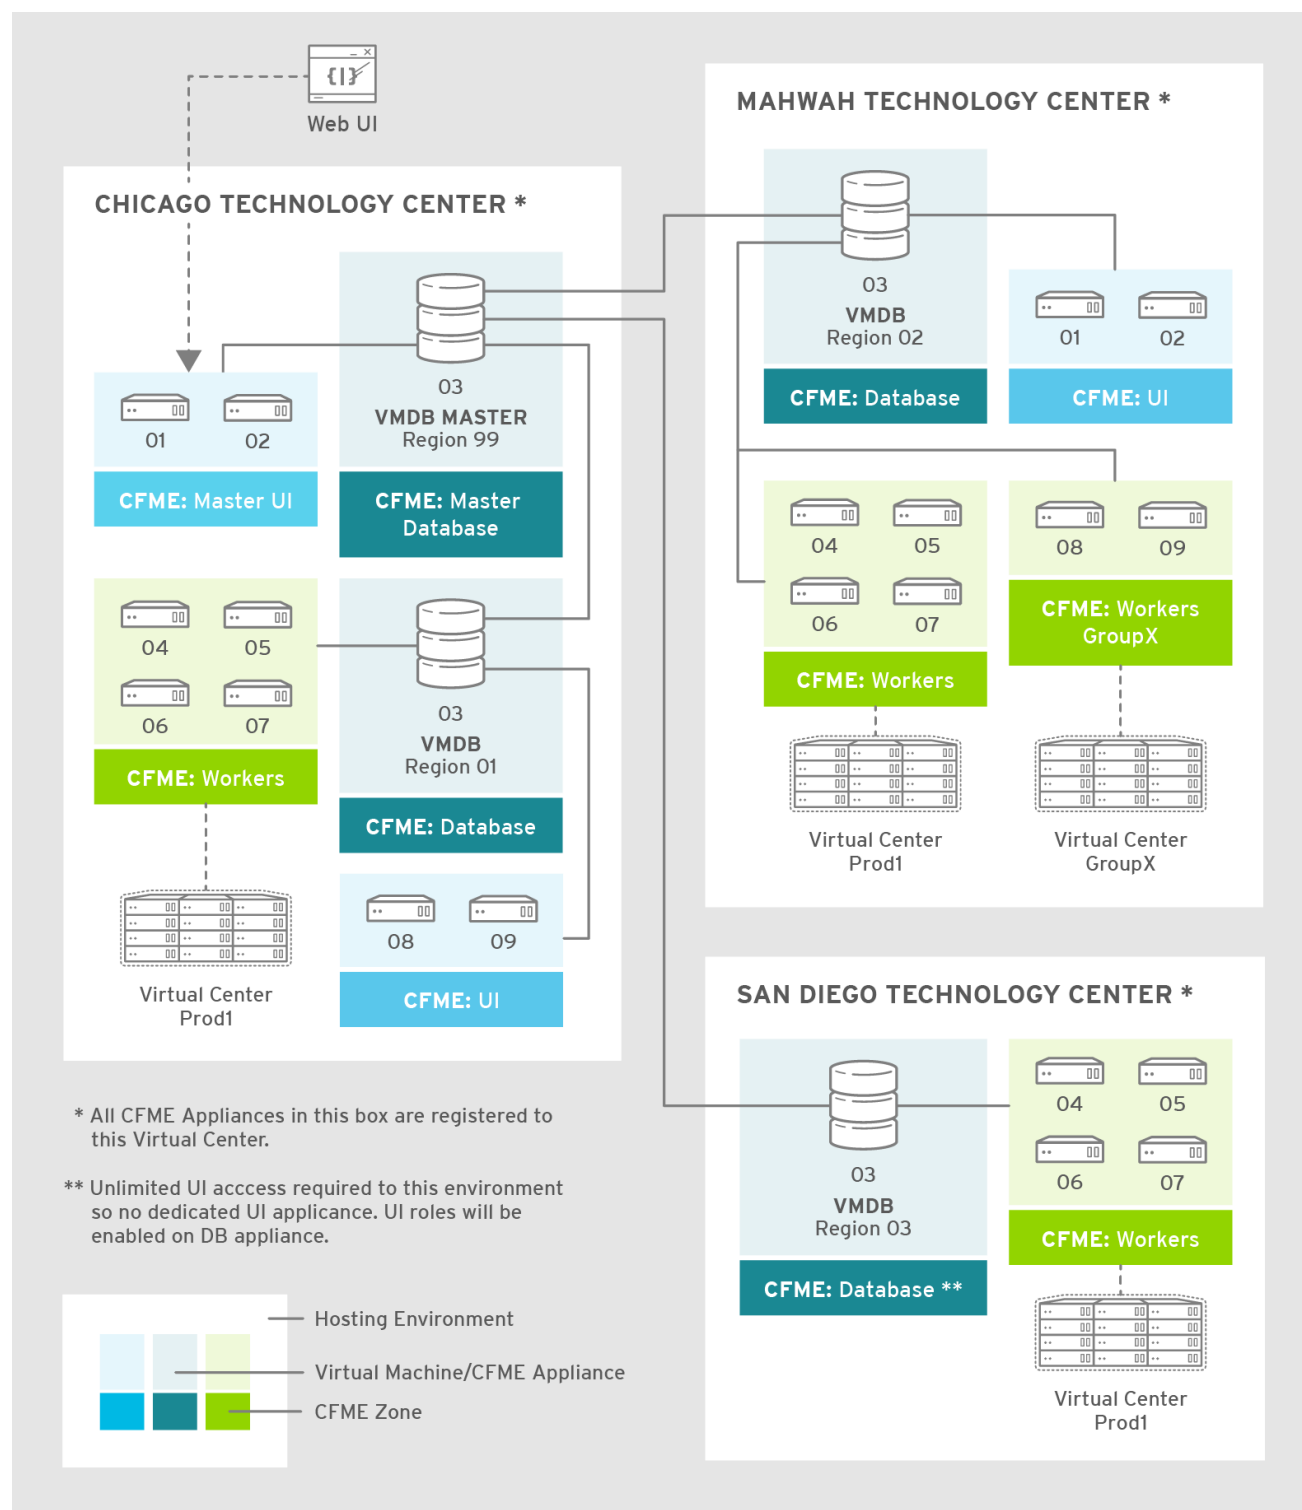
Use **Regions** for centralizing data which is collected from public and private virtualization environments. A region is ultimately represented as a single database for the VMDB. Regions are particularly useful when multiple geographical locations need to be managed as they enable all the data collection 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 CloudForms environment.



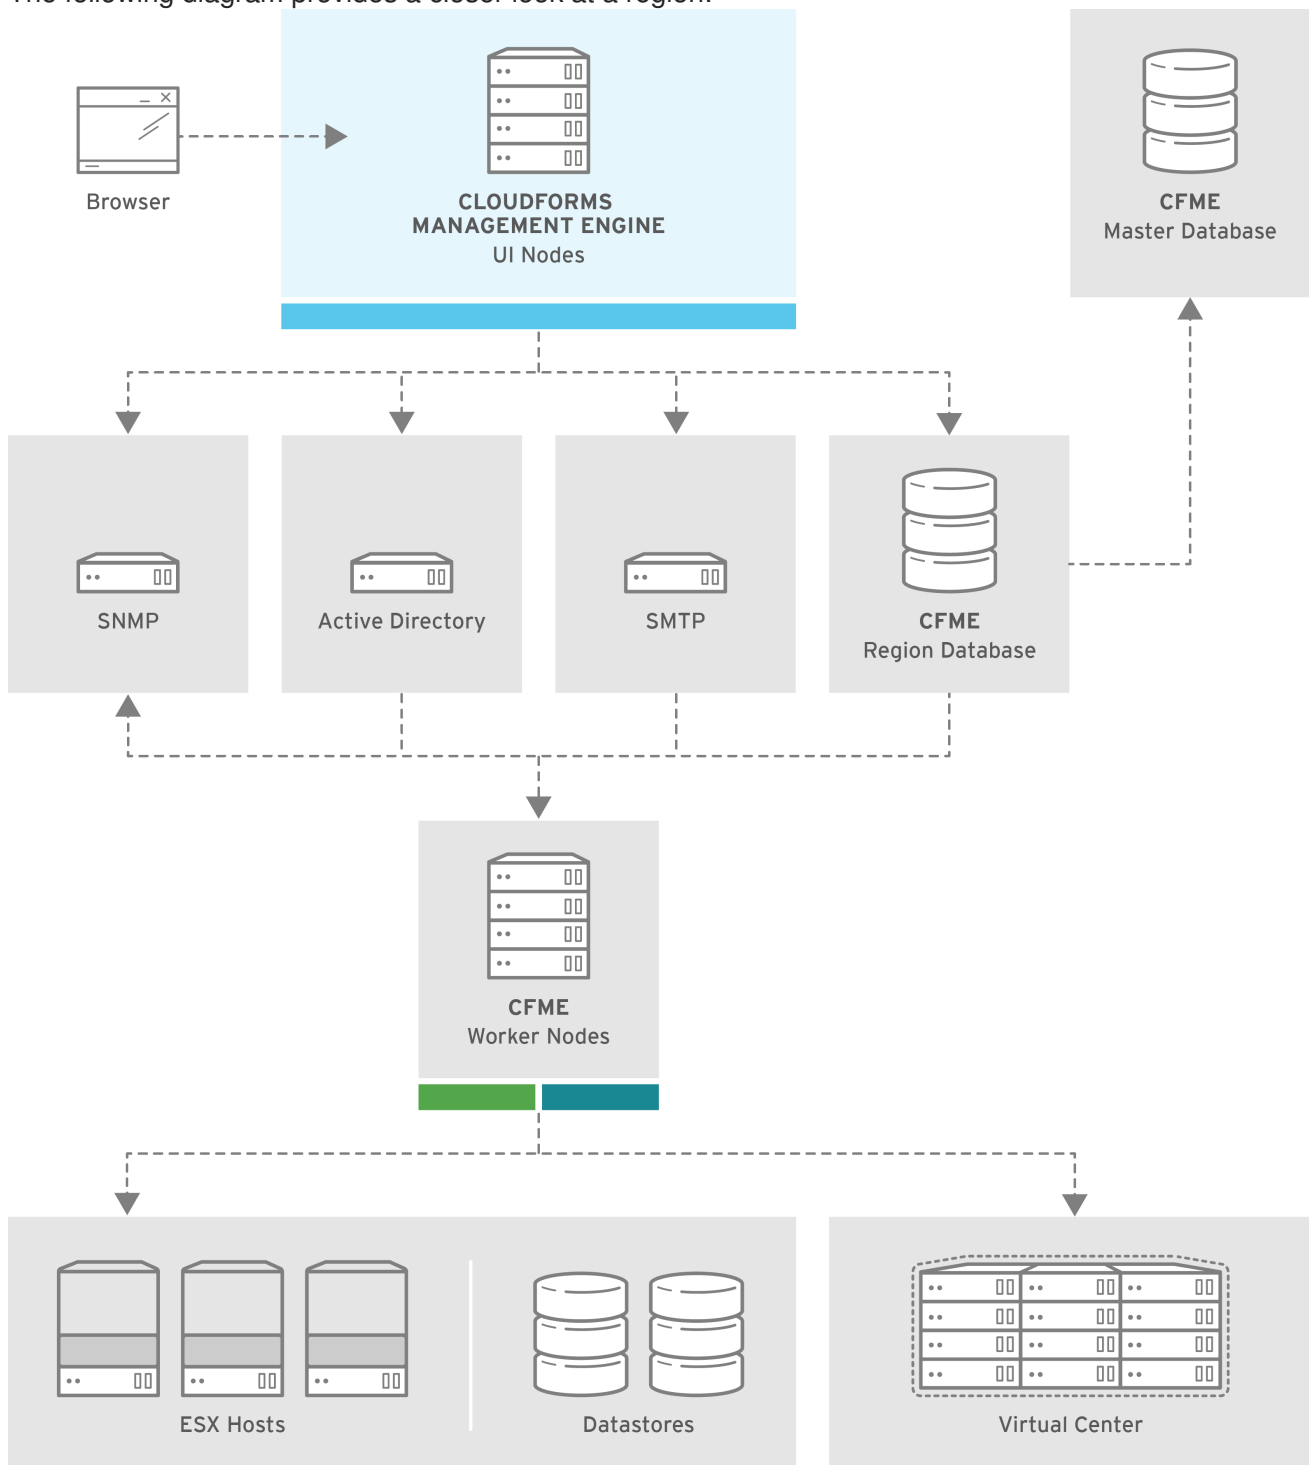
The Master appliance is located in Chicago and contains a master region and a subregion that manages the worker appliances. The Mahwah technology center contains a single subregion that manages two zones. Likewise the San Diego technology center contains a single subregion managing a single zone.



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).

3.1.1.1. 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.

3.1.1.2. Region Settings

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

3.1.1.3. Capacity and Utilization Collections

3.1.1.3.1. 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, the clusters, hosts, and datastores show under **Configure** → **Configuration**, then by clicking on the **Settings** → **Region** → **C & U Collection** tab.

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

To Enable a Cluster, Host, or Datastore for Capacity and Utilization Collection:

1. Navigate to **Configure** → **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

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.

3.1.1.4. Tags

3.1.1.4.1. Company Tag Categories and Tags

CloudForms Management Engine 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 by navigating to **Configure** → **Configuration**, then clicking on the **Settings**, then selecting **Region**, then the **My Company Tags** tab. 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 CloudForms Management Engine.
- **System tags** are assigned automatically by CloudForms Management Engine.




NOTE

If you entered a Company Name under **Configure** → **Configuration**, then clicking on the **Settings** tab, then the **Server** your desired server, that name will appear on the tab instead of **My Company**.

3.1.1.4.2. Creating a Tag Category

To Create a Tag Category:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Categories** tab.

3. Click  (Click to add a new category).
4. In the **Category Information** area,

Category Information

Name	<input type="text"/>
Display Name	<input type="text"/>
Description	<input type="text"/>
Show In Console	<input checked="" type="checkbox"/>
Single Value	<input checked="" type="checkbox"/>
Capture C & U Data by Tag	<input type="checkbox"/>

* '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 **Display Name** to specify how you want to see the name of the category in the Console.
- 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.
- Check **Show** in Console when you feel that the category is ready for use in the console. For example, you want to populate values for the category before exposing it to users.
- Check **Single Value** 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.
- Check **Capture C & U Data** by tag to be able 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.

3.1.1.4.3. Deleting a Tag Category

To Delete a Tag Category:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Categories** tab.
3. Click  (**Click to delete this category**) next to the category to delete it.



NOTE

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

3.1.1.4.4. Creating a Company Tag

To Create a Company Tag:

1. Navigate to **Configure** → **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  (**New Entry**), and type a **Name** and **Description** for your new value.


5. Click  (**Add this entry**) to confirm the entry.


6. Repeat these steps for each value you need.

3.1.1.4.5. Deleting a Company Tag

To Delete a Company Tag:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then **Region**, then click on the **My Company Tags** tab.

3. Click  (**Click to delete this entry**) 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.

3.1.1.4.6. 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
```

3.1.1.4.7. 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. Navigate to **Configure** → **Configuration**.
3. Click on the **Settings** accordion, then **Region**, then click on the **Import Tags** tab.
4. Click **Browse** 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**.

3.1.1.4.8. 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.

3.1.1.4.9. 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. Navigate to **Configure** → **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 **Browse** 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**.

3.1.1.5. Registering and Updating Red Hat CloudForms

You can register appliances, edit customer information, and update appliances from the **Red Hat Updates** tab, accessible from **Configure** → **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 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 CloudForms content every 12 hours.



NOTE

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

3.1.1.5.1. Registering Appliances

Before you can access and apply package updates, you must register and subscribe the 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

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

Red Hat Software Updates

The Subscription Management Service you register with will provide your systems with updates and allow additional management.

Register to: **Red Hat Subscription Management**

Red Hat Subscription Management Address: **Red Hat Satellite 6**

Repository Name(s): **cf-me-5.5-for-rhel-7-rpms rhel-server-rhsc1-7-rpms**

HTTP Proxy: ☐ Use HTTP Proxy

Enter your Red Hat account information

Login:

Change Password / Confirm Password:

Validate

Forgot your login or password? Look it up at http://redhat.com/forgot_password

Save **Reset** **Cancel**

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. Navigate to **Configure** → **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 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-5.5-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-5.5-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**.

Appliance Updates

Available Product version: 5.5.4.2

<input type="checkbox"/>	Appliance	Zone	Update Status	Last Checked for updates	CFME Version	Platform Updates Available	Last Message
<input type="checkbox"/>	EVM	default	Subscribed	2016-06-20 02:15:36 UTC	5.5.4.2	Yes	Registered

You can now apply updates to your appliance.

3.1.1.5.2. Updating Appliances

An important part of securing Red Hat CloudForms is to ensure your appliances use the latest packages. Package updates to the appliance contain patches for any software bugs, including possible security bugs.

The **Red Hat Updates** tab enables you to check for updates and update registered appliances. Any services requiring a restart to apply updates are automatically restarted as part of the **Red Hat Updates** process.

1. Navigate to **Configure** → **Configuration**.
2. Select **Region** in the accordion menu and click the **Red Hat Updates** tab.
3. Click **Check For Updates** to search the Content Delivery Network (CDN) for any updated CloudForms packages. If an appliance update is available, it will be listed with the available version.
4. Click **Apply CFME Update** to install and update CloudForms packages. The CloudForms service will be automatically restarted as needed.



IMPORTANT

Red Hat recommends updating the appliance using the **Red Hat Updates** tab only, which automatically restarts services and the appliance if needed. Running updates from the command line requires manually restarting the CloudForms service and sometimes the appliance. If updating the CloudForms appliance from the command line, use the **yum -y update cfme-appliance** command to update only the CloudForms appliance and its dependencies, and to avoid any potential incompatibilities that could be introduced by running **yum update**.

The following options are available in the **Appliance Updates** section of **Red Hat Updates**:

Table 3.1. Options Available in Appliance Updates

Option	Use
Refresh List	Refreshes the list of appliances.
Check for Updates	Checks for all available CloudForms updates using yum .
Register	Attempts to register the appliance if it is not already registered. CloudForms subscribes to the rhel-x86_64-server-6-cf-me-3 RHN channel for RHN-registered appliances, and to the products designated by Red Hat product certification for subscription-manager registered appliances. The Red Hat Enterprise Linux channels are enabled by default on registration. In addition, CloudForms automatically checks for updates after registering.
Apply CFME Update	Applies updates to CloudForms packages only. Specifically, this option runs the yum -y update cfme-appliance command. This command installs every package listed in the dependency tree if it is not already installed. If a specific version of a package is required, that version of the package is installed or upgraded. No other packages, such as PostgreSQL or Red Hat Enterprise Linux, are updated. The appliance may be rebooted automatically during this process.



NOTE

If the appliance is registered to Red Hat Satellite, you can use content views to manage updates for CloudForms. For more information, see [Creating Content Views](#) in the *Red Hat Satellite 6 Content Management Guide*.

3.1.1.5.3. Updating the Appliance Operating System

Updating the appliance's operating system requires a manual update using the **yum** command. This command updates all RPMs on the appliance, not just CloudForms packages, and is supported only when used as part of migration.



IMPORTANT

Scheduled downtime is required while updating system packages for the following reasons:

- Some updates may interrupt CloudForms operations.
- Updates for the PostgreSQL database server suspend CloudForms operations.
- System updates may require a reboot of the CloudForms appliance.

Red Hat recommends updating the appliance using the **Red Hat Updates** tab in the CloudForms user interface, which automatically restarts services and the appliance if needed. See [Section 3.1.1.5.2, “Updating Appliances”](#) for instructions.

To update the appliance operating system:

1. Log in to the appliance console as the root user.
2. Run the **yum update** command and confirm any updates.
3. Restart the appliance as required.

3.1.1.5.4. 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.

For more information on managing subscriptions for an IT infrastructure, see the Red Hat Subscription Management [Subscription Concepts and Workflows](#) guide.

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.

For more information on migrating from older Red Hat Network Classic (hosted) to the updated subscription management, see the Red Hat Subscription Management [Migrating from RHN Classic](#) guide.



3.1.2. Profiles


3.1.2.1. 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.


3.1.2.2. Creating a Host Analysis Profile


To Create a Host Analysis Profile:


1. Navigate to **Configure** → **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	<input type="text"/>
Description	<input type="text"/>
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 CloudForms Management Engine Control. See the *Defining Policies and Profiles* guide, available from <https://access.redhat.com/documentation/en/red-hat-cloudforms/>.
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	<input type="text"/>	warn	<input type="text"/>	5

9. Click **Add**.

3.1.2.3. Creating a Virtual Machine Analysis Profile

To Create a Virtual Machine Analysis Profile:


1. Navigate to **Configure** → **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	<input type="text"/>
Description	<input type="text"/>
Type	Vm


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




Category Selection


☐ Services
 ☐ Software
 ☐ System
 ☐ User Accounts
 ☐ VM Configuration

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


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



File Entry

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


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

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

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

11. From the **Event Log Entry** area, click  (**Click to add a new entry**) 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.




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**.

3.1.2.4. Editing an Analysis Profile


To Edit an Analysis Profile:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then click **Analysis Profiles**.
3. Check the analysis profile you want to edit.
4. Click  (**Edit this Analysis Profile**).
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.

3.1.2.5. Copying an Analysis Profile

To Copy an Analysis Profile:


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

3.1.2.6. 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. Navigate to **Configure** → **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 this 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**.

3.1.3. Zones

You can organize your CloudForms Management Engine 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 CloudForms Management Engine 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 CloudForms Management Engine 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 CloudForms Management Engine Appliances in the CloudForms Management Engine 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. Default zone is provided. This can be removed only after you have created your own zones.

3.1.3.1. Creating a Zone

To Create a Zone:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click  (**Configuration**), and  (**Add a new zone**) to create a zone.

4. In the **Zone Information** area, type in a **Name** and **Description** for the new zone.

Zone Information	
Name	West
Description	Western Zone
SmartProxy Server IP	192.168.252.12

5. 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.
6. 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	
User ID	acme\admin
Password
Verify Password

7. 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.
8. In the **Settings** area, set the number for **Max Active VM Scans**. The default is Unlimited.
9. Click **Save**.

3.1.3.2. Deleting a Zone



To Delete a Zone:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. 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.

3.1.3.3. Editing a Zone

To Edit a Zone:

1. Navigate to **Configure** → **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**.

3.1.4. Servers

Server settings enables you to control how each CloudForms Management Engine 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.

3.1.4.1. Changing Server Settings

To Change Server Settings:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Settings** accordion, then click **Zones**.
3. Click the zone where the CloudForms Management Engine server is located.
4. In the **Servers** area, click on the CloudForms Management Engine server.
5. Click **Server**.
6. Make any required changes.
7. Click **Save**.

3.1.4.1.1. Basic Information Settings

Basic Information

Hostname	localhost.localdomain.localdomain
IP Address	10.64.15.217
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 CloudForms Management Engine 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 **Configure** → **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.

3.1.4.1.2. Server Control Settings

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 CloudForms Management Engine appliance, be sure to set this on all of the appliances.

3.1.4.1.3. Server Roles


Server Role	Description
Automation Engine	Use this role if you want to use this CloudForms Management Engine server to process automation tasks.
Capacity and Utilization (three server roles)	<ul style="list-style-type: none"> • 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 Management Engine server in a specific zone has this role, only one will be active at a time. • 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 Management Engine server with this role in a zone. • The Capacity & Utilization Data Processor processes all of the data collected, allowing CloudForms Management Engine to create charts. This role has a dedicated worker, and there can be more than one CloudForms Management Engine server with this role in a zone.
Database Operations	Use Database Operations to enable this CloudForms Management Engine server to run database backups or garbage collection.
Database Synchronization	Use Database Synchronization to enable this CloudForms Management Engine server's VMDB to replicate to a higher-level VMDB. This should only be enabled after creating settings for the Replication Worker.

Server Role	Description
Event Monitor	This role is enabled by default and provides the information shown in timelines. Event Monitor is responsible for the work between the CloudForms Management Engine 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 Management Engine 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.
Notifier	Use this role if you will be using CloudForms Management Engine Control or Automate to forward SNMP traps to a monitoring system or send e-mails. See Section 3.1.4.1.6, “Configuring SNMP” for details on creating SNMP alerts. If more than one CloudForms Management Engine server in a specific zone has this role, only one will be active at a time.
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 Management Engine server in a specific zone has this role, only one will be active at a time.
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 Management Engine Control. More than one CloudForms Management Engine server can have this role in a zone.
RHN Mirror	An appliance with RHN Mirror enabled acts as a server containing a repository with the latest CloudForms Management Engine 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.
Reporting	This role is enabled by default. The Reporting role specifies which CloudForms Management Engine servers can generate reports. If you do not have a CloudForms Management Engine 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.
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 zone must be assigned this role or scheduled CloudForms Management Engine events will not occur. If more than one CloudForms Management Engine server in a specific zone has this role, only one will be active at a time.

Server Role	Description
SmartProxy	Enabling the SmartProxy role turns on the embedded SmartProxy on the CloudForms Management Engine 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 Management Engine console. This SmartProxy can also analyze virtual machines on the host on which it is installed.
SmartState Analysis	This role is enabled by default. The SmartState Analysis role controls which CloudForms Management Engine servers can control SmartState Analyses and process the data from the analysis. You should have at least one of these in each zone.
User Interface	This role is enabled by default. Uncheck User Interface if you do not want users to be able to access this CloudForms Management Engine server using the CloudForms Management Engine console. For example, you may want to turn this off if the CloudForms Management Engine server is strictly being used for capacity and utilization or reporting generation. More than one CloudForms Management Engine server can have this role in a zone.
Web Services	This role is enabled by default. Uncheck Web Services to stop this CloudForms Management Engine server from acting as a Web service provider. More than one CloudForms Management Engine server can have this role in a zone.

3.1.4.1.4. VMware Console Settings

If you are using the CloudForms Management Engine 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, the appropriate version of the VMware MKS plug-in or the appropriate VMRC viewer installed in your Web browser.



NOTE

- You are responsible for installing the correct version for your virtual infrastructure. See vendors documentation for information. After installing the appropriate software or version, you must specify which version you are using in the CloudForms Management Engine configuration settings.
- To edit the VMware MKS plug-in settings, you must have the super administrator role.

The screenshot shows a configuration window titled "VMware Console Support". It contains three labeled input areas: "Use" with a dropdown menu currently showing "VNC", "VNC Proxy Address" with an empty text box, and "VNC Proxy Port" with an empty text box.

- 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 MKS** plug-in, select the appropriate version.
- If using **VMware VMRC** plug-in, be sure that you have fulfilled the requirements. The correct version of the VMRC plug-in from VMware must be installed on the client computer. To do this, log into the Virtual Center Web Service and attempt to open a virtual machine console. This should prompt you to install the required plug-in. The VSphere Web Client must be installed on VC version 5, and the provider must be registered to it. For Virtual Center version 4, the VMware VirtualCenter Management Webservice must be running.

3.1.4.1.5. NTP Servers Settings

In the **NTP Server** 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.

3.1.4.1.6. 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 CFME 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. Navigate to **Configure** → **Configuration** → **Settings**.
3. Select the zone where the EVM server is located, and select the EVM server.
4. In the **Server Control** area, select the **Notifier** server role.
5. Click **Save**.

To Set the Failover Priority Role:

1. Navigate to **Configure** → **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.

3.1.4.1.7. Outgoing SMTP Email Settings

To use the email action in CloudForms Management Engine, 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"/>
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 domain name for the mail server.
- Check **Start TLS Automatically** if the mail server requires TLS.
- Select the appropriate verify mode.
- Use the **Authentication** drop down to specify if you want to use login or plain authentication.
- Use **User Name** to specify the user name required for login authentication.
- Use **Password** to specify the password for login authentication.
- Use **From Email Address** to set the address you want to send the email from.
- Use **Test Email Address** if you want to test your email settings. Click **Verify**.

To Test Outgoing SMTP Email Server Settings:

1. Type in all settings in the Outgoing SMTP Email Server settings, including Test Email Address.
2. Click  (Send test email).

3.1.4.1.8. Web Services Settings

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




Web Services

Mode	invoke
Security	none

- Set **Mode** to **invoke** to enable 2-way Web services communication between the CloudForms Management Engine Appliance and the SmartProxy. Set **Mode** to **disabled** to use Web services from the SmartProxy to the CloudForms Management Engine Appliance only. When the CloudForms Management Engine 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 CloudForms Management Engine 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 **ws-security**.

3.1.4.1.9. Logging Settings




Logging

Log Level	info
------------------	------

- 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*.

3.1.4.1.10. Custom Support URL Settings



Custom Support URL

URL (i.e. www.mypage.com)	www.redhat.com
Description	Red Hat

- 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**.

3.1.4.2. 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.

3.1.4.2.1. Changing an Authentication Setting

To Change an Authentication Setting:

1. Navigate to **Configure** → **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 on the **Authentication** tab.
6. Make any required changes. If you select LDAP, LDAPS, or Amazon as the authentication mode, click Validate to confirm your settings in the Role Settings area.
7. Click **Save**.

3.1.4.2.2. Authentication Settings

The screenshot shows the 'Authentication' settings panel. It has a title 'Authentication' at the top. Below it, there are two main sections. The first section is 'Session Timeout', which has a text input field containing '1' followed by a unit selector 'h' and another text input field containing '0' followed by a unit selector 'm'. The second section is 'Mode', which has a dropdown menu currently showing 'Database'.

- Use **Session Timeout** to set the period of inactivity before a user is logged out of the console.
- Use **Mode** to set the type of authentication. Choose from **Database** (using the VMDB), **LDAP** (Lightweight Directory Authentication Protocol), **LDAPS** (Secure Lightweight Directory Authentication Protocol), or **Amazon**. The default is **Database**. If you choose **Database**, see Section **Creating a User** to create users. See section **LDAP Settings** for more information on configuration for LDAP and LDAPS. If you choose Amazon, see section **Amazon Settings**.

3.1.4.2.3. LDAP Settings

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

LDAP Settings

The screenshot shows the 'LDAP Settings' form. It has several input fields. The first section is 'LDAP Host Names', which has three stacked text input fields. The second section is 'LDAP Port', which has a text input field containing '389'. The third section is 'User Type', which has a dropdown menu currently showing 'User Principal Name'. The fourth section is 'Domain Prefix:', which has a text input field containing '<domain>\<user>'. The fifth section is 'User Suffix: <user>@', which has a text input field.

- Use **LDAP Host Name** to specify the fully qualified domain names of your LDAP servers. CloudForms Management Engine will search each host name in order until it finds one that authenticates the user.
- 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 **User Principal Name** to type the user name in the format of user@domainname. Select **Email Address** to login with the users email address. Select **Distinguished Name** (CN=<user>) or **Distinguished Name** (UID=<user>) to use just the user name, but be sure to enter the proper **User Suffix** for either one. Choose the correct **Distinguished Name** option for your directory service implementation.
- 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**.

3.1.4.2.4. LDAP Role Settings

If you choose LDAP, you can use groups from your directory service to set the role for the authenticated LDAP User. The LDAP user must be in one of the Account Role Groups. See Section “LDAP Groups”.

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. For example, dbright@acme.com when using User Principal Name, cn=dan bright,ou=users,dc=acme,dc=com when using Distinguished Name (CN=<user>), or uid=dan bright,ou=users,dc=acme,dc=com when using Distinguished Name (UID=<user>). Then, when logging in, the user would type either dbright (User Principal Name) or dan bright (Distinguished Name). If the user is not defined in the VMDB, they will be denied access to CloudForms Management Engine.

Role Settings

Get User Groups from LDAP	<input checked="" type="checkbox"/>
Get Roles from Home Forest	<input type="checkbox"/>
Follow Referrals	<input type="checkbox"/>
Base DN	dc=acme,dc=com
Bind DN	dbright@acme.com
Bind Password

Validate

- Check **Get Roles** from Home Forest to use the LDAP roles from the LDAP users home forest.
- Check **Follow Referrals** to lookup and bind a user that exists in a domain other than the one configured in the LDAP authentication settings.

- Use **Base DN** to set the place in the directory tree from which you want to start searching for users.
- 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 Management Engine authentication and role assignment.
- Specify the password for the Bind DN user in **Bind Password**.

Click **Validate** to verify your settings.

3.1.4.2.5. 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 Management Engine 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 Management Engine user or belong to an IAM user group added to the list of CloudForms Management Engine groups.

3.1.4.2.6. Trusted Forests


If a user has group memberships in another LDAP Forest, then 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 Management Engine 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 Management Engine. The user must be a member of at least one matching LDAP group to be successfully authenticated.

3.1.4.2.7. Adding Settings for a Trusted Forest

To Add Settings for a Trusted Forest:

1. Navigate to **Configure** → **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 **Authentication**.
6. Check **Get Role** 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**.

3.1.4.2.8. External Authentication (httpd)

When external authentication is enabled, users can log in to the CloudForms Management Engine 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 **Configure** → **Configuration** → **Access Control** → **Groups** to administer groups.

Configuring 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 **Configure** → **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 **Configure** → **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 **Configure** → **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 3.1. 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 3.2. Reverting to Internal Database Authentication

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

3.1.4.3. 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.

3.1.4.3.1. Changing Settings for a Worker

To Change the Settings for a Worker (except replication worker)

1. Navigate to **Configure** → **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**.

3.1.4.3.2. Changing Settings for the Replication Worker




IMPORTANT

This should only be entered on subordinate servers that will have the Database Synchronization role enabled. These settings must be completed before enabling that role.

To Change Settings for the Replication Worker:

1. Navigate to **Configure** → **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 **Replication Worker** area.



Replication Worker	
Database	vmdb_production
Port	5432
Username	root
Password
Verify Password
Host	
<input type="button" value="Validate"/>	

- Use **Database** 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.
 - Use **Host** to specify the IP address or hostname of the top level VMDB.
7. Click **Validate** to confirm that the VMDB is accessible.
8. Click **Save**.

The new settings take one to two minutes to take effect. Next, you need to enable the replication worker on the subordinate regions VMDB server.

3.1.4.4. 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.

3.1.4.4.1. Changing a Database Setting

To Change a Database Setting:

1. Navigate to **Configure** → **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.

3.1.4.5. Customization and Logos

3.1.4.5.1. Custom Logos

Use Custom Logos to display your own logo in the corner of the console or on the CloudForms Management Engine login panel.


3.1.4.5.2. Uploading a Custom Logo to the Console

To Upload a Custom Logo to the Console:

1. Make sure the desired logo is accessible from the computer where you are running the console. The file must be in portable network graphics (png) format with dimensions of 350 x 70.
2. Navigate to **Configure** → **Configuration**.
3. Click on the **Settings** accordion, then click **Zones**.
4. Click the zone where the CloudForms Management Engine server is located.
5. Click on the server.

- Click the **Custom Logos** tab.

Custom Logo Image (Shown on top right of all screens)




- Click **Browse** in the **Custom Logo Image** (Shown on top right of all screens) area to go to the location where the logo file is located.
- Click **Upload**. The icon is displayed above the file name box, and an option is shown to use the logo.
- Check **Use Custom Logo** Image to add the logo to your console.
- Click **Save**.

3.1.4.5.3. Customizing the Login Background

To Customize the Login Background:

- Make sure the logo that you want to use is accessible from the computer where you are running the console. The file must be in PNG format with dimensions of 1280 x 1000.
- Navigate to **Configure** → **Configuration**.
- Click on the **Settings** accordion, then click **Zones**.
- Click the zone where the server is located.
- Click on the server.
- Click the **Custom Logos** tab.
- Click **Browse** in the **Custom Login Background Image** area to go to the location where the

Custom Login Background Image



logo file is located.

- Click **Upload**. The icon is displayed above the file name box, and an option is shown to use the logo.
- Check **Use Custom Login Image** to add the logo to your console.
- Click **Save**.

3.1.4.5.4. Customizing the Login Panel Text

To Customize the Login Panel Text:

1. Navigate to **Configure** → **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 Panel Text**, type in text that you want to show on the consoles login screen.



Custom Login Panel Text (0 / 500)

Use Custom Login Text ☐

7. Check **Use Custom Login Text** box to add the text to the screen.
8. Click **Save**.

3.1.4.6. 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 CloudForms Management Engine server.



NOTE

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

3.1.4.6.1. Editing Configuration Files Manually

To Edit Configuration Files Manually:

1. Navigate to **Configure** → **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**.

3.1.4.6.2. Configuration Parameters

Table: authentication

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 CFME 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 CFME 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 CFME 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 CFME authentication and role assignment. This is the same as Bind DN in Configuration-Operations-Server-LDAP Settings in the CFME 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 CFME 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 CFME 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 CFME 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 CFME 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 CFME 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 CFME 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 CFME 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, CFME 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
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 CFME 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. CFME 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

Parameters	Description
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
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 CFME. 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 CFME 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 CFME 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
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
------------	-------------

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 CFME Appliance log. This is the same as Log Path in Configuration-Operations- Server-Logging in the CFME 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

Parameters	Description
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	<p>Specifies what is considered archived logging. The default pattern is blank which means *.gz files in the log directory.</p>

Table 3.9. log_depot

Parameters	Description
uri	<p>Specify the uri for the log depot. This is the same as URI in Configure → Configuration → Diagnostics Collect Logs in the CFME Console. Default: blank</p>
username	<p>Specify the user name for the log depot. This is the same as User ID in Configure → Configuration → Diagnostics Collect Logs in the CFME Console. Default: blank</p>
password	<p>Specify the password for the user for the log depot. This is the same as Password in Configure → Configuration → Diagnostics Collect Logs in the CFME Console. Default: blank</p>

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	
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

Parameters	Description
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 CFME 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
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 CFME 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 CFME Appliance. Default: 2.minutes
host	CFME Server's IP address. Default: blank
hostname	CFME Server's hostname. Default: localhost.localdomain

Parameters	Description
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 CFME 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 CFME Appliance that you are logged on to in the CFME Console. This is the same as Appliance Name in Configuration-Operations- Server-Basic Information. Default : EVM
role	Specify the roles for this CFME Server, separated by commas without spaces. The possible values are automate, database_operations, database_synchronization, 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 CFME 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

Parameters	Description
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
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 CFME 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 CFME 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 CFME 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 CFME 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 CFME 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

Parameters	Description
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
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

Parameters	Description
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.
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 CFME 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

Parameters	Description
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
replication_worker:	Performs database replication tasks. Settings for Database Synchronization Server Role. Parameters specified here will override those set in the queue_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: 200.megabytes
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
replication:	This section contains information for the destination database for the replication.
destination:	
database	Name of destination database. Default: vmdb_production
username: root	Username for the destination database. Default: root
password	Stores password for destination database in encrypted format.
host	Host of the destination database.
port	Port of the destination database. Default: 5432
include_tables	Lists tables included in the replication. Do NOT modify unless specifically instructed to do so by ManageIQ support. Default: include all, exclude_tables is used instead.
exclude_tables	Lists tables not to be included in the replication. Do NOT modify unless specifically instructed to do so by ManageIQ support.
options	
replication_trace	Set to true to capture all replication tracing in the log. Default: false

Parameters	Description
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
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 CFME 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 CFME Server statistics. Default: 5.minutes
server_stats_interval	How often to collect the CFME Server statistics. Default: 60.seconds
session_timeout_interval	How often to check to see if a UI (CFME Console) session has timed out. Default: 30.seconds
storage_file_collection_interval	How often to perform file inventory of storage locations. Default: 1.days

Parameters	Description
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
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

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: 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
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

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: 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
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

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: 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
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

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: 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
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



3.1.5. Schedules

3.1.5.1. 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.

3.1.5.1.1. Scheduling a SmartState Analysis or Compliance Check

To Schedule a SmartState Analysis or Compliance Check:

1. Navigate to **Configure** → **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



Basic Information	
Name	<input type="text"/>
Description	<input type="text"/>
Active	<input checked="" type="checkbox"/>
Action	<div>VM Analysis ▼</div> <div> VM Analysis Template Analysis Host Analysis Cluster / Deployment Role Analysis Datastore Analysis VM Compliance Check Host Compliance Check Database Backup </div>

- **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**.
- By applying **Global Filters** within any of the above items, you can designate which virtual machines or hosts to analyze.
 - In the **Timer area**, click the **Run** list to set the frequency of the analysis to run. There are further options based on which **Run** option you choose.

Timer	
Run	Once ▾
Time Zone	(GMT+00:00) UTC ▾ below
Starting Date	3/7/2013
Starting Time (UTC)	0 ▾ h 0 ▾ m

- 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.

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

3.1.5.2. Scheduling a Database Backup

To Schedule a Database Backup:

1. Navigate to **Configure** → **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.

Basic Information	
Name	DB daily backup
Description	DB daily backup
Active	<input checked="" type="checkbox"/>
Action	Database Backup ▾

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**.

Basic Information	
Name	DB daily backup
Description	DB daily backup
Active	<input checked="" type="checkbox"/>
Action	Database Backup ▾

- If selecting **Samba**, enter the **Depot Name**, **URI**, **User ID**, and a valid **Password**. Then, click **Validate** to check the settings.

validate to check the settings.

- If you choose **Network File System**, enter the **Depot Name** and **URI**.
8. In the **Timer** area, click the **Run** list to specify how often to run the analysis. Your options after that depend on the **Run** option you choose.

Timer

Run	Daily  every Day 
Time Zone	(GMT+00:00) UTC 
Starting Date	7/30/2013
Starting Time (UTC)	0  h 0  m

- 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 analysis.
 - Click **Hourly** to run the backup hourly. You will be prompted to select the number of hours between each analysis.
9. Select a **Time Zone**.





NOTE

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

3.1.5.2.1. Modifying a Schedule

To Modify a Schedule:



1. Navigate to **Configure** → **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**.

3.2. ACCESS CONTROL

From navigating to menu: Configure[Configuration], then clicking on the Access Control accordion, you have a hierarchy of the configurable items for users, groups, roles, and tenants. You can add and modify users, groups, account roles, and tenants.



3.2.1. Creating a Tenant

To Create a Tenant:

1. Navigate to **Configure** → **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**.



3.2.2. Creating a Project

To Create a Project:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Access Control** accordion, then click **Tenants**.
3. Click on the **Tenant** where you want to add a **Project**, click  (**Configuration**), and  (**Add Project to this Tenant**) to create a Project.
4. Enter a name for the Project in the **Name** field.
5. Enter a description for the Project in the **Description** field.
6. Click **Add**.

3.2.3. Creating a Project Quota

To Create a Project Quota:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Access Control** accordion, then click **Tenants**.
3. Click on the **Project** where you want to add a Quota, click  (**Configuration**), and  (**Manage Quotas**) to create a Quota.
4. In the list of pre-built Quotas, mark the **Enforced checkbox** next to the Quota item you want to enable.
5. In the **Value** field, enter the constraints you want to apply to the Quota.

6. Click **Save**.



3.2.4. Tagging Tenants and Projects

To tag Tenants and Projects:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Access Control** accordion, then click **Tenants**.
3. In the **Tenant or Project** entry, click the **Policy** drop-down menu, and select **Edit My Company Tags for this Tenant**.
4. In the **Tag Assignment** table, 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**.

3.2.5. Creating a User

To Create a User:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Access Control** accordion, then click **Users**.
3. Click  (**Configuration**), and  (**Add a new User**) to create a user.
4. Type in a **Full Name**, **Username**, **Password** with confirmation, **Email Address**, and choose a **Group** for the user.

Adding a new User

User Information

Full Name	<input type="text"/>		
Username	<input type="text"/>		
Change Password / Confirm Password	<input type="text"/>	<input type="text"/>	
E-mail Address	<input type="text"/>		
Group	<Choose a Group> ▼		



NOTE

- 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 to be sure that the user can be found in the VMDB. For example, `jdunn@acme.com` when using User Principal Name, `cn=Jack Dunn,ou=users,dc=acme,dc=com` when using Distinguished Name (CN=<user>), or `uid=JackDunn,ou=users,dc=acme,dc=com` when using Distinguished Name (UID=<user>).
- Then, when logging in, the user would type either `jdunn` for User Principal Name or `Jack Dunn` for Distinguished Name. If the user is not defined in the VMDB, they will be denied access to CloudForms Management Engine. The password field will not be used.
- When the user logs in they should use their LDAP password.



5. Select a Group.

6. Click Add.

3.2.6. 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. Navigate to **Configure** → **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.

3.2.7. Groups

User groups create filters and assign roles to users. You can either create your own user 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 "Roles".

3.2.8. Creating a User Group

To Create a User Group:

1. Navigate to **Configure** → **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 **Description** 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 **Tenant/Project** 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 a bold, blue 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 a bold, blue font.

Assign Filters

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

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

▶ ☐ ab.usersys.redhat.com

☐ AmazonDan

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 a bold, blue font.

Assign Filters

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

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

▶ ☐ ab.usersys.redhat.com

☐ AmazonDan

13. Click **Add**.

3.2.9. LDAP Groups

When leveraging your LDAP groups, if you are using LDAP and the LDAP user is not a member of any of the defined groups, then the user will be denied access to CloudForms Management Engine. There are two ways to use LDAP groups with CloudForms Management Engine:

- Create groups with a specific set of names as provided by CloudForms Management Engine. These groups automatically get assigned to a specific role.
- Assign pre-existing groups from your LDAP server to CloudForms Management Engine account roles.

3.2.10. Using CloudForms Management Engines Named Groups to Assign Account Roles

In your directory service, define a distribution group for each 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 "LDAP Settings".

3.2.11. Account Role and Directory Service Group Names



Directory Service Distribution Group Name	Account Role
EvmGroup-administrator	Administrator
EvmGroup-approver	Approver
EvmGroup-auditor	Auditor
EvmGroup-desktop	Desktop
EvmGroup-operator	Operator
EvmGroup-security	Security
EvmGroup-super_administrator	Super Administrator
EvmGroup-support	Support
EvmGroup-user	User
EvmGroup-user_limited_self_server	User Limited Self Service
EvmGroup-user_self_service	User Self Service
EvmGroup-vm_user	Vm User
EvmRole-tenant_administrator	Tenant Administrator
EvmRole-tenant_quota_administrator	Tenant Quota Administrator


1. Make each user of your directory service that you want to have access to CloudForms Management Engine a member of one of these groups.

2. Navigate to **Configure** → **Configuration**, then click on the **Settings** accordion, then **Zones**, then the **Authentication** tab, you can enable **Get User Groups from LDAP** after typing in all of the required settings. See Section **LDAP Settings**.

3.2.12. Using Pre-existing LDAP Groups to Assign Account Roles

To Use Pre-existing LDAP Groups to Assign Account Roles:


1. Navigate to **Configure** → **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 description for the group in the **Description** field.
5. There are two ways to specify the group to use:
 - Type in the **cn** for the group in **LDAP Group**. This group must be in the LDAP directory source you specified under Operations-Server.



Group Information

Description	HelpDesk	<input checked="" type="checkbox"/> (Look Up LDAP Groups)
Role	EvmRole-administrator ▼	

- Select **Look Up LDAP Groups** to find a list of groups, then use the drop down list that appears in the **LDAP Group Information** area to choose a group.




LDAP Group Look Up

User to Look Up	<input type="text"/>
Username	<input type="text"/>
Password	<input type="password"/>

Retrieve

6. Select a **Role** to map to the group.
7. Select any **filters** to apply to what this group can view in the **Assign Filters** area:
 - a. Select the filters to apply to the user. The items that have changed show in a bold, blue font.
 - b. Click the **Host & Clusters** tab.


- c. Select the host and clusters to limit the user to. The items that have changed show in a bold, blue font.




Assign Filters

My Company Tags Hosts & Clusters VMs & Templates

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

▶ ☒  **ab.usersys.redhat.com**

☐  AmazonDan

- d. Click the **VMs & Templates** tab. This shows folders that you have created in your virtual infrastructure.
- e. Select the folders to limit the user to. The items that have changed show in a bold, blue font.
8. Click **Add**.

3.2.13. Roles

When you create a user group, you must specify a role to give the group rights to resources in the console, and then assign a user to a group. CloudForms Management Engine provides a default group of roles, but you can also create your own as well as copy the default groups. The table below shows the function available to each group.



NOTE

If you have enabled Get Role from LDAP under LDAP Settings, then the role is determined by the LDAP users group membership in the directory service. See Section "LDAP Settings".

3.2.13.1. Account Roles and Descriptions



Role	Description
Administrator	Administrator of the virtual infrastructure. Can access all infrastructure functionality. Cannot change server configuration.
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.
Desktop	Access to VDI pages.


Role	Description
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 CloudForms Management Engine 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.
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.

3.2.14. Creating a Role

To Create a Role:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Access Control** accordion, then click **Roles**.

- Click  (Configuration), and  (Add a new Role).
- 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 specifically used by the user, by the user or its group, or all virtual machines.



Role Information


Name: EvmRole-user

VM & Template Access Restriction: None

Options: None, Only User or Group Owned, Only User Owned

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

Product Features (Editing)



- Click **Add**.

3.3. DIAGNOSTICS

From navigating to **Configure** → **Configuration**, then clicking on the **Diagnostics** tab, you can also see the status of the different CloudForms Management Engine 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 CloudForms Management Engine 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.

3.3.1. 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 CloudForms Management Engine 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.

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**, **Database Synchronization**, **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.





3.3.2. Region Aware Server Roles

Role	More than one per Region	Can have Priority Set
Automation Engine	Y	N

Role	More than one per Region	Can have Priority Set
Database Operations	Y	N
Database Synchronization	N	Y
Notifier	N	Y
Reporting	Y	N
Scheduler	N	Y
User Interface	Y	N
Web Services	Y	N

3.3.3. Setting the Priority of a Failover Role

To Set the Priority of a Failover Role:

1. Navigate to **Configure** → **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.

3.3.4. Replication

You must be on the server where replication has been set up to monitor status. To run backups, the database operations server role must be enabled. Databases can then be restored using the black console on the CloudForms Management Engine Appliance. These features are available only when using the internal PostgreSQL VMDB.

3.3.4.1. Monitoring Database Replication

To Monitor Database Replication:

1. Navigate to **Configure** → **Configuration**.

2. Click the **Diagnostics** accordion and click the **Region** being replicated.
3. Click the **Replication** tab to view the status of the replication process and the backlog of database records queued for replication.

If directed to by Red Hat, you may need to reset replication. Do this from the server that is replicating up to a higher level VMDB. When you do this, the subordinate regions data is removed from the top level, and then the replication is restarted.

3.3.4.2. Resetting Replication

To Reset Replication:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Diagnostics** accordion, then click Region name.
3. Click the **Replication** tab.
4. Click **Reset**.

3.3.4.3. Running a Single Backup

To Run a Single Backup:

1. Navigate to **Configure** → **Configuration**.
2. Click on the **Diagnostics** accordion, then click Region name.
3. Click the **Database** tab.
4. If you have created a backup schedule, and want to use the same depot settings, select it under **Backup Schedules**.
5. If you do not want to use the settings from a backup schedule, or need to create settings, go into the **Database Backup Settings** area.
6. Select a type of server to put the backups. You can either use **Network File System** or **Samba**.

Database Backup Settings

Type	Samba
URI	smb:// backupserver.acme.com/backups
User ID	acme/admin
Password
Verify Password

Validate

- If selecting **Samba**, enter the **Depot Name**, **URI**, **User ID**, a **Password**, and a valid **Password**. Click **Validate** to check the settings.
- If you choose **Network File System**, enter the **Depot Name** and **URI**.

7. Click **Submit**.

The database backup is run immediately. You can see the task by navigating to **Configure** → **Tasks**, then clicking on the **All Other Tasks** tab.

3.3.4.4. Restoring a Database from Backup

To Restore a Database from Backup:

1. Copy the database backup file to **/tmp**, and name it **evm_db.backup**. The server looks specifically for this file to restore from.
2. Log in to the black appliance console with a user name of **admin** and the default password. The CloudForms Management Engine Appliance summary screen displays.
3. Press **Enter** to manually configure settings.
4. Press **6** to select Restore Database From Backup.
5. Press **Y** to confirm.

If directed by Red Hat, you can run database garbage collection to reclaim unused space in your VMDB. Generally, the database server does this automatically.

3.3.5. 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.

3.3.5.1. Viewing the Status of Server Roles

To View the Status of Server Roles:

1. Navigate to **Configure** → **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**.

3.3.5.2. Setting 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 are **Notifier**, **Capacity & Utilization Coordinator**, **Database Synchronization**, **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 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.
- **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 are 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.



3.3.5.3. 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
Database Synchronization	N	Y
Event Monitor	N	Y
Provider Inventory	N	Y
Provider Operations	Y	N
Notifier	N	Y
Reporting	Y	N
Scheduler	N	Y
SmartProxy	Y	N

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


3.3.5.4. Setting the Priority of a Failover Role

To Set the Priority of a Failover Role:

1. Navigate to **Configure** → **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. From 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  (**Promote Server to primary for this role**) to make this the primary Server for this role.
6. Click  (**Demote Server to normal for this role**) to demote the priority of this Server for this role.

3.3.5.4.1. Removing an Inactive Server

To Remove an Inactive Server:

1. Navigate to **Configure** → **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.

3.3.5.5. Zone Log Collections

Zone Log Collection Settings

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



- Anonymous File Transfer Protocol (FTP)


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

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

3.3.5.5.1. Setting the Location of the Log Depot

To Set the Location of the Log Depot:

1. Navigate to **Configure** → **Configuration**.
2. Click the **Diagnostics** accordion, then click the **Zone** that you want to view.
3. Click **Collect Logs**.
4. Click  (**Edit the Log Depot Settings for the selected Zone**).
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 <No Depot>


- <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**.

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

To Collect and Download Logs from all Servers in a Zone:

1. Navigate to **Configure** → **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 logs**). All files in the logs directory as well as configuration files are collected.
5. Click **OK**. The status of the log retrieval shows in the CloudForms Management Engine console.


3.3.5.6. Capacity and Utilization Repair

Under certain circumstances, it is possible that CloudForms Management Engine 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 CloudForms Management Engine 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.

3.3.5.6.1. Repairing Capacity and Utilization Data

To Repair Capacity and Utilization Data

1. Log in to a CloudForms Management Engine Appliance located in the zone for which you want to gather the data.
2. Navigate to **Configure** → **Configuration**.
3. Click on the **Diagnostics** accordion, then click the **Zone** that you want to view.
4. Click **C & U Gap Collection**.



- Select the appropriate **Timezone**.



NOTE

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

- Select a **Start Date**.
- 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.

3.3.6. Server Diagnostics

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

3.3.6.1. Workers

The Workers tab enables you to see the status of and restart CloudForms Management Engine 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.
- **Database Replication Worker** that is responsible for maintaining replication activities.
- **Event Handlers** put events from the Event Monitor into the VMDB and starts CloudForms Management Engine 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 CloudForms Management Engine Web services.
- **VM Analysis Collectors** that run and process SmartState Analyses on virtual machines.

3.3.6.1.1. Reloading Worker Display



To Reload Worker Display:

1. Navigate to **Configure** → **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**).

3.3.6.1.2. Restarting a CloudForms Management Engine Worker

To Restart a CloudForms Management Engine Worker:

1. Navigate to **Configure** → **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**.

3.3.6.2. Server and Audit Logs

3.3.6.2.1. 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


- 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.

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

To Set the Location of the Log Depot for a Specific Server

1. Navigate to **Configure** → **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	<div><No Depot></div> <div> <div><No Depot></div> <div>Anonymous FTP</div> <div>FTP</div> <div>NFS</div> <div>Red Hat Dropbox</div> <div>Samba</div> </div>
------	---

7. 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	FTP
Depot Name	
URI	ftp://
User ID	
Password	
Verify Password	
<div>Validate</div>	

8. Enter your user ID and password, then click **Validate** to confirm the settings.
9. Click **Save**.

3.3.6.2.3. Collecting the Current Log Set of a Server



To Collect the Current Log Set of a Server

1. Navigate to **Configure** → **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 CloudForms Management Engine console.

3.3.6.2.4. Collecting All Log Sets from a Server

To Collect All Log Sets from a Server

1. Navigate to **Configure** → **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 CloudForms Management Engine console.

3.3.6.2.5. 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 **Configure** → **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.

3.3.6.2.6. Viewing the Server Log

To View the Server Log


1. Navigate to **Configure** → **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 CloudForms Management Engine server automatically retrieves the last 1000 lines of the log.

3.3.6.2.7. Reloading the Server Log


To Reload the Server Log

1. Navigate to **Configure** → **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).

3.3.6.2.8. Downloading the Server Log

To Download the Server Log

1. Navigate to **Configure** → **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 Log File**).



NOTE

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

3.3.6.2.9. Viewing the Audit Log


To View the Audit Log:

1. Navigate to **Configure** → **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.


3.3.6.2.10. Reloading the Audit Log

To Reload the Audit Log

1. Navigate to **Configure** → **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**).

3.3.6.2.11. Downloading the Audit Log

To Download the Audit Log

1. Navigate to **Configure** → **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).

3.3.6.2.12. Viewing the Production Log

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


To View the Production Log

1. Navigate to **Configure** → **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 CloudForms Management Engine server automatically retrieves the last 1000 lines of the log.

3.3.6.2.13. Reloading the Production Log

To Reload the Production Log

1. Navigate to **Configure** → **Configuration**.
2. Click on the Diagnostics accordion, then click the Zone that you want to view.
3. Click **Production Log**.
4. Click the **CloudForms Management Engine Log** tab.
5. Click  (**Reload the Product Log Display**).

3.3.6.2.14. Downloading the Production Log

To Download the Production Log

1. Navigate to **Configure** → **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**).

3.4. DATABASE OPERATIONS

3.4.1. 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. Navigate to **Configure** → **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.

3.4.2. Database Regions and Replication

3.4.2.1. 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.



NOTE

Only enable database synchronization on subordinate servers with replication worker settings already configured. Do not enable more than one Database Synchronization role per region.

For diagnostic information on replication, see Section "Monitoring Database Replication".

The following is an example of regionalized database scenario:

1. Create **Region Number 99** to which all other VMDBs replicate.
 - Treat this as a **read only** database for reporting and charting.


- Enable only the **Reporting**, **Scheduler**, and **User Interface Server** Roles. To perform database maintenance items, such as scheduled backups, on the top-level region (master), also enable the **Database Operations** role.
- No additional settings aside from assigning the region ID. No need to configure any replication.

2. Create Region Number 1

- Add replication worker settings pointing to the VMDB for **Region 99**.
- Enable **Database Synchronization Server** role on one **Server** in the **Region**. If you have a second **Server** in the same region, do not enable the **DB Synchronization** role. Do not enable more than one **Database Synchronization Role** per Region.

3. Create Region Number 2

- Add replication worker settings pointing to the VMDB for **Region 99**.
- Enable **Database Synchronization Server** role on one **Server** in the **Region**. If you have a second **Server** in the same region, do not enable the **DB Synchronization** role. Do not enable more than one **Database Synchronization Role** per Region.



IMPORTANT

All CFME databases in a multi-region deployment must use the same security key.

3.4.2.2. Creating a Region

Use this procedure to create a region on a CloudForms Management Engine (CFME) Appliance that already has a database. The process of creating a region takes a few minutes. The database is dropped and rebuilt to accommodate the region numbers.



NOTE

In most cases, a region is created when you are deploying your CFME environment; essentially, on the first Appliance for the region. To configure a database for CFME, follow the procedure described [here](#).

To Create a Region:

1. Start the appliance and log in to the black appliance console with a user name of admin and the default password. The **Appliance Summary Screen** displays.
2. Press **Enter** to manually configure settings.
3. Enter **13** to Stop EVM Server Processes.
4. Enter **Y** to confirm.
5. The menu reappears after all processes are complete.
6. Choose **Configure Database**.

**WARNING**

Performing this action destroys any existing data and cannot be undone. Back up the existing database before proceeding. By default, new CloudForms Management Engine Appliances are assigned region 0. Do not use this number when creating a region as duplicating region numbers can compromise the integrity of the data.

7. Enter a database region number that has not been used in your environment. Do not enter duplicate region numbers as this can corrupt the data.
8. Enter **Y** to confirm.
9. The menu reappears after all processes are complete.
10. Enter **14** to Start EVM Server Processes.
11. Enter **Y** to confirm.

After a region is created, you can create subordinate regions as necessary and set up replication to the top level region.

3.4.3. Replicating a Database

To Replicate a Database

1. Navigate to **Configure** → **Configuration**.
2. Click the **Settings** accordion and click **Zones**.
3. Click the **Zone** where the server is located and click the server name.
4. Click **Workers**.
5. In the **Replication Worker** area, enter the worker information:
 - a. **Database**: the name of your VMDB.
 - b. **Port**: the Port number on which the VMDB service is running.
 - c. **Username**: the user name to connect to the VMDB user name.
 - d. **Password** and **Verify Password**: the password for the user name.
 - e. **Host**: the IP address or hostname of the top level VMDB.
6. Click **Validate** to confirm that the VMDB is accessible.
7. Click **Save**.

3.4.4. Enabling the Database Synchronization Role

When you start the replication worker, all of the information in the subordinate database is sent to the top region (99). The worker also creates triggers so that future changes made to subordinate regions are sent automatically to the top region.

To Enable the Database Synchronization Role:

1. Navigate to **Configure** → **Configuration**.
2. Click the **Settings** accordion and click **Zones**.
3. Click the **Zone** where the server is located and click the server name.



NOTE



Only enable database synchronization on subordinate servers with replication worker settings already configured. Do not enable more than one Database Synchronization role per region.

4. Click Server.
5. In the Server Control area, select Database Synchronization.
6. Click Save.

3.4.4.1. Scheduling a Database Backup

Schedule database backups on a regular basis to preserve data.

To Schedule a Database Backup

1. Navigate to **Configure** → **Configuration**.
2. Click the **Settings** accordion and click **Schedules**.
3. Click  (**Configuration**), and  (**Add a new Schedule**).
4. In the **Basic Information** box, enter a **Name** and **Description** for the schedule.

Basic Information	
Name	DB daily backup
Description	DB daily backup
Active	<input checked="" type="checkbox"/>
Action	Database Backup ▾

5. Check **Active** to enable the backup schedule.
6. In the **Action** drop-down list, select **Database Backup**.

7. In the **Database Backup Settings** box, select a type of server for storing the backups from the **Type** drop-down list. You can use **Network File System** (NFS) or **Samba**.

Database Backup Selection

Type	Samba
Depot Name	BackupServer
URI	smb:// backupserver.acme.com/backups
Username	acme/admin
Password	••••••
Verify Password	••••••
Validate	

- 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**.
8. In the **Timer** box, select the desired backup frequency from the **Run** list.

Timer

Run	Daily every Day
Time Zone	(GMT+00:00) UTC
Starting Date	7/30/2013
Starting Time (UTC)	0 h 0 m

- **Once**: the backup runs one time.
 - **Hourly**: select the number of hours between backups from the drop-down list.
 - **Daily**: select the number of days between backups from the drop-down list.
 - **Weekly**: select the number of weeks between backups from the drop-down list.
 - **Monthly**: select the number of months between backups from the drop-down list.
9. Select a **Time Zone** for the schedule.

10. Type or select a **Starting Date** for the schedule.
11. Select a **Starting Time** based on a 24 hour clock.
12. Click **Add**.

3.4.4.2. Running a Single Database Backup

To Run a Single Database Backup

1. Navigate to **Configure** → **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**.

Database Backup Selection

Type	Samba
Depot Name	BackupServer
URI	smb:// backupserver.acme.com/backups
Username	acme/admin
Password
Verify Password
<input type="button" value="Validate"/>	

- 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.

3.4.4.3. Restoring a Database from a Backup

If a database is corrupt or fails, restore it from a backup.

To Restore a Database from a Backup

1. Save the database backup file as **/tmp/evm_db.backup**. CloudForms Management Engine looks specifically for this file when restoring a database from a backup.
2. Log in to the black appliance console with a user name of admin and the default password. The **Appliance Summary Screen** displays.
3. Press **Enter** to manually configure settings.
4. Enter **13** to Stop Server Processes. Stop the process on all servers that connect to this VMDB.
5. Enter **Y** to confirm.
6. After all processes are stopped, press **Enter** to return to the menu.
7. Press **Enter** again to manually configure settings.
8. Enter **6** to select Restore Database From Backup. If connections are open, the server may still be shutting down. Wait a minute and try again.
9. Enter **y** to keep the database backup after restoring from it. Enter **n** to delete it.
10. Press **Y** to confirm.
11. After the backup completes, press **Enter** to return to the menu.
12. Press **Enter** again to manually configure settings.
13. Enter **14** to Start Server Processes.
14. Enter **Y** to confirm.

3.4.4.4. 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. Navigate to **Configure → 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**.

3.4.5. Changing Database Password

3.4.5.1. Changing the Password on the Database Appliance

CloudForms Management Engine provides a default database password for the internal PostgreSQL database.

To change the password, you need to stop the CloudForms Management Engine Service, change the password for the PostgreSQL database, run a command to change the password in the configuration file that emserverd uses to access the server, and restart the CloudForms Management Engine Appliance.

1. Stop the CloudForms Management Engine Service.

- **SSH** into the appliance.
- To stop the CloudForms Management Engine service, run the following command:

```
service evmserverd stop
```

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

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

- At the vmdb# prompt, type:

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

- To exit psql, type:

```
\q
```

3. Run the following command to change the password in the configuration file that emserverd uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseym1 --password newpassword
```

4. To restart the CloudForms Management Engine service, run the following command:

```
service evmserverd start
```

5. Verify that you can log in to the CloudForms Management Engine Console.

3.4.5.2. Changing the Password on the Worker Appliances

1. Stop the CloudForms Management Engine Service.

- **SSH** into the appliance.
- To stop the CloudForms Management Engine service, run the following command:

```
service evmserverd stop
```

2. Run the following command to change the password in the configuration file that emserverd uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseym1 --password newpassword
```

- To restart the CloudForms Management Engine service, run the following command:

```
service evmserved start
```

IMPORTANT

Requisites for Replication

If you are using replication, and you have changed the password on the **Master** database, you will also need to update the password used for the **Replication Worker** on the subordinate regions.

- Disable the **Database Synchronization** Role.
- Update the **Replication Worker's** password.
- Restart the **Database Synchronization** Role.

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

- Create the new appliance.
- Start it, but do not go into any of the configuration options, instead **SSH** into the new appliance.
- 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.)
- Edit the `database.yml` file in the `/var/www/miq/vmdb` directory. (You may want to save off the original.)
 - Replace the contents of the **"production"** section with the contents of the **"base"** section.
 - Edit the **"host"** parameter to match IP of the appliance hosting the VMDB.
 - Save the new `database.yml`.
 - Run the following command to change the password in the configuration file that emserved uses to access the server:

```
/var/www/miq/vmdb/tools/fix_auth.rb --databaseym1 --password  
newpassword
```

- Restart the new worker appliance by typing:

```
service evmserved restart
```

CHAPTER 4. SMARTPROXIES

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

4.1. 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. See the *Managing Infrastructure and Inventory* guide, available from <https://access.redhat.com/documentation/en/red-hat-cloudforms/>, for information on discovery.

4.2. 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 **Infrastructure** → **Hosts**.
2. Select the host you want to edit.
3. Click  (**Configuration**), then  (**Edit this item**).
4. In the Credentials box, 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

User ID	root
Password
Verify Password

Validate

Required. Should have privileged access, such as root or administrator.



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.

1. Click **Validate** to verify the credentials.
2. 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.
3. 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 located by navigating to **Configure** → **Configuration**, then clicking on the desired server, then the **Server** tab, and exploring the **Server Control** area.

CHAPTER 5. ABOUT

In the **About** area of **Configure**, you can see session information for the console, download the Red Hat CloudForms documentation in PDF format, and navigate to the Red Hat Customer Portal.


To view session information, documentation, and the Red Hat Customer Portal link:

1. Navigate to **Configure** → **About**.
2. To go to the Red Hat Customer Portal, click access.redhat.com.
3. To view the documentation, click the title of the document to view.

CHAPTER 6. RED HAT ACCESS INSIGHTS

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

Red Hat Access 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 Access 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 Access 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 Access Insights plugin has the following options:

- Overview
- Rules
- Systems


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

6.1. 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.



6.2. 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.

6.2.1. 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



6.2.2. 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.

6.3. 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

<input type="checkbox"/>	Hostname	Last Check In	Status
<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	