Red Hat CloudForms 5.0-Beta

Monitoring, Alerts, and Reporting

Creating and managing reports, feeds, and widgets in Red Hat CloudForms

Red Hat CloudForms 5.0-Beta Monitoring, Alerts, and Reporting

Creating and managing reports, feeds, and widgets in Red Hat CloudForms

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

Legal Notice

Abstract

This guide provides instructions for creating and managing reports, feeds, and widgets in Red Hat CloudForms. It also includes information on accessing usage and chargeback costs. This information supports better information technology decision making and predictions for future virtual machine management. 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. OVERVIEW DASHBOARD

Overview shows your virtual environment's events, reports, and configurable alerts. This information supports better information technology decision making and predictions for future virtual machine management.

When you log in to the console, it brings you directly to the **Overview Dashboard** page by default. The console uses widgets to organize this page, providing you with a default set of commonly used widgets. The configuration items on these widgets are clickable, allowing you to drill down directly to a referenced item.

Customize this page to include charts and reports you specifically want to see as soon as you log in to the console. You can add, remove, move, minimize, zoom into, and maximize widgets. Only users with the proper access can create widgets.

1.1. ADDING A WIDGET

You can add widgets to the dashboard to accommodate the information you want to see upon login.

1. Navigate to **Overview → Dashboard**.

2. Click ✚ **(Add a Widget)**.

3. Select the widget you want to add from the list.

> **NOTE**
>
> Only widgets that are not currently showing on the dashboard will appear in this list.

1.2. RESETTING TO THE DEFAULT SET OF WIDGETS

You can reset to default set of widgets according to your needs with the following steps:

1. Navigate to **Overview → Dashboard**.

2. Click ↩ **(Reset Dashboard Widgets to the defaults)**

1.3. REMOVING A WIDGET

You can remove a widget when you no longer need the widget to be displayed in the dashboard.

1. Navigate to **Overview → Dashboard**.

2. From the widget that you want to remove, click ✖ **(Remove from Dashboard)** in the upper right corner of the widget.

3. Click **OK**.

The widget is removed from the dashboard, but it is not deleted. It can be added again if needed.

1.4. ZOOMING IN TO A CHART WIDGET

1. Navigate to **Overview → Dashboard**.

2. From the chart widget that you want to enlarge, click ✚ **(Zoom in)** in the upper right corner of the widget.

The enlarged chart is opened in its own pop-up window. To close this window, click ✖ **(Close)** in the upper right corner.

1.5. OPENING A CHART OR REPORT WIDGET IN FULL SCREEN VIEW

1. Navigate to **Overview → Dashboard**.

2. From the chart or report widget that you want to enlarge, click ⛶ **(Full Screen view)** in the upper right corner of the widget.

3. Click **OK**.

1.6. MINIMIZING OR MAXIMIZING A WIDGET

1. Navigate to **Overview → Dashboard**.

2. From the chart or report widget that you want to maximize, click ▣ **(Minimize)** or ⛶ **(Full Screen)** in the upper right corner of the widget.

1.7. DOWNLOADING A REPORT WIDGET AS A FILE

1. Navigate to **Overview → Dashboard**.

2. From the chart or report widget that you want to download as a PDF, click 📄 **[(Download PDF)** in the upper right corner of the widget.

3. Click **OK**.

CHAPTER 2. CREATING DASHBOARD WIDGETS

2.1. CREATING A REPORT WIDGET

1. Navigate to **Overview → Reports**.

2. Click on the **Dashboard Widgets** accordion, then choose the **Reports** folder.

3. Click ⚙ (**Configuration**), then click ⊕ (**Add a new Widget**).

4. In the **Basic Information** area, type in a **Title** and **Description**. By default the widget will be active as soon as you create it. To make it inactive, uncheck the **Active** box.

5. From the **Report Options** area, select the filters until you are at the report you want for this widget. Select up to four columns from that report. Finally, for **Row Count**, select the number of rows that you want displayed.

6. In the **Timer** area, click the **Run** drop down to specify how often you want the widget data to get updated. The options displayed will depend on which **Run** option you choose. Select **Hourly**, **Daily**, **Weekly**, or **Monthly**.

7. Select a **Time Zone**. Type or select a date to begin the schedule in **Starting Date**. Select a **Starting Time** based on a 24 hour clock in the selected time zone.

> **NOTE**
>
> If you change the time zone, you will need to reset the starting date and time.

8. In the **Visibility** area, select **<To All Users>**, so that all users can use this widget no matter what user role they are assigned. Select **<By Role>** to assign this widget to specific user roles. Select **<By Group>** to assign this widget to specific groups.

9. Click **Add**.

2.2. CREATING A CHART WIDGET

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widgets** accordion and click the **Charts** folder.

3. Click ⚙ (**Configuration**), then click ⊕ (**Add a new Widget**).

4. In the **Basic Information** area, type in a **Title** and **Description**. By default the widget will be active as soon as you create it. To make it inactive, uncheck the **Active** box.

5. From the **Chart Report** area, select a chart to display in the widget.

Chart Report	
Filter	Configuration Management/Virtual Machines/VMs with Free Space > 50% by Department ⬍

6. In the **Timer** area, click the **Run** drop down to specify how often you want the widget data to get updated. The options displayed will depend on which **Run** option you choose. Select **Hourly**, **Daily**, **Weekly**, or **Monthly**.

Timer	
Run	Hourly ⬍ every Hour ⬍
Time Zone	(GMT+00:00) UTC ⬍ fields below
Starting Date	3/14/2013
Starting Time (UTC)	0 ⬍ h 0 ⬍ m

7. Select a **Time Zone**. Type or select a date to begin the schedule in **Starting Date**. Select a **Starting Time** based on a 24 hour clock in the selected time zone.
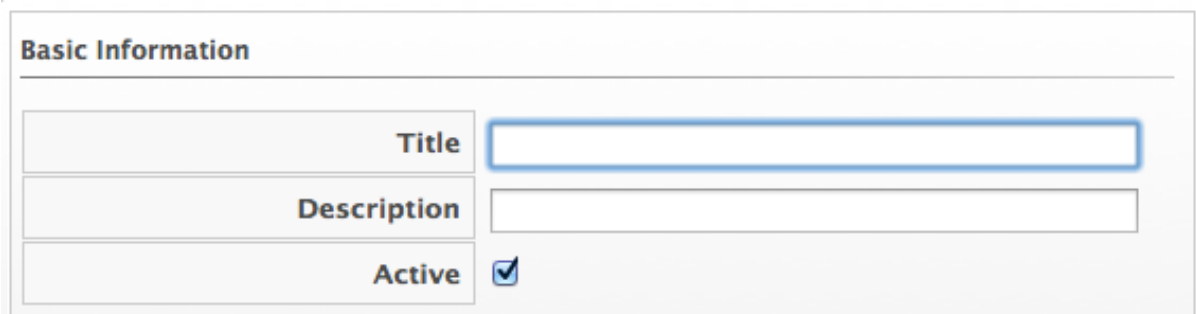
> **NOTE**
>
> If you change the time zone, you will need to reset the starting date and time.

8. In the **Visibility** area, select **<To All Users>**, so that all users can use this widget no matter what user role they are assigned. Select **<By Role>** to assign this widget to specific user roles. Select **<By Group>** to assign this widget to specific groups.

9. Click **Add**.

2.3. CREATING A MENU WIDGET

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widgets** accordion and click the **Menus** folder.

3. Click ⚙ (**Configuration**), then click ⊕ (**Add a new Widget**).

4. In the **Basic Information** area, type in a **Title** and **Description**. By default the widget will be active as soon as you create it. To make it inactive, uncheck the **Active** box.

5. In the **Menu Shortcuts** area, use the **Add a Shortcut** dropdown to select all the places in the console that you want to add to this widget.

6. In the **Visibility** area, select **<To All Users>**, so that all users can use this widget no matter what user role they are assigned. Select **<By Role>** to assign this widget to specific user roles. Select **<By Group>** to assign this widget to specific groups.

7. Click **Add**.

2.4. EDITING A WIDGET

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widgets** accordion and select the widget you want to edit.

3. Click ⚙ (**Configuration**), and then ✏ (**Edit this Widget**).

4. Make the required changes.

5. Click **Save**.

2.5. COPYING A WIDGET

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widget** accordion, and select the widget you want to copy.

3. Click ⚙ (Configuration), and then ▤ (Copy this Widget).

4. Type a unique name for the widget and edit its properties.

5. Click **Save**.

2.6. DELETING A WIDGET

> **NOTE**
>
> Default widgets cannot be deleted, but they can be copied.

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widget** accordion and select the widget you want to delete.

3. Click ⚙ (Configuration), and then 🗑 (Delete this Widget from the Database)

4. Click **OK**.

2.7. IMPORTING A WIDGET

You can share widgets between appliances using the export and import features.

1. Navigate to **Overview → Reports**.

2. In the **Import/Export** accordion, click **Widgets**.

3. In the **Import** area, click **Browse** to select an import file.

4. Click **Upload**.

2.8. EXPORTING A WIDGET

You can share widgets between appliances using the export and import features.

1. Navigate to **Overview → Reports**.

2. In the **Import/Export** accordion, click **Widgets**.

3. In the **Export** area, select the widgets that you want to export.

4. Click **Export**.

2.9. GENERATING WIDGET CONTENT IMMEDIATELY

1. Navigate to **Overview → Reports**.

2. Click the **Dashboard Widgets** accordion and select the widget you want to generate.

3. Click ⚙ (Configuration), and then ⚙ (Generate Widget Content now).

4. Click **OK**.

The content is generated immediately instead of waiting for the next scheduled update. Generation of widget content is shown under **Tasks**, accessed from the settings menu.

CHAPTER 3. REPORTS

Click the **Reports** accordion under **Overview → Reports** to see a list of reports available. These reports have been constructed to help you view the most commonly requested and significant data. From here, you can also create reports if you have appropriate access. Red Hat CloudForms provides a large group of default reports organized into categories. Each category has its own set of subfolders.

- **Configuration Management**: hardware, application, network, service, user account, operating system, and snapshot information for all of your items.

- **Migration Readiness**: information related to items required to migrate a virtual machine.

- **Operations**: free space on registered and unregistered virtual machines, power states for virtual machines, and SmartState analysis status.
  This category also provides reports relating to the operation of Red Hat CloudForms, such as user IDs and snapshots taken by Red Hat CloudForms.

- **VM Sprawl**: usage information and disk waste.

- **Relationships**: virtual machine, folder, and cluster relationships.

- **Events**: operations and configuration management events.

- **Performance by Asset Type**: performance of your virtual infrastructure.
  You must be capturing capacity and utilization data to get this information.

- **Running Processes**: information on processes running on a virtual machine.
  You must have domain credentials entered for the zone to collect the information for these reports, and the virtual machine must have been analyzed at least once.

- **Trending**: projections of datastore capacity, along with host CPU and memory use.

- **Tenants**: quotas report aggregated by each tenant that shows quota name, total quota, in use, allocated, and available. The report currently lists all tenants and there is no nesting information available by parent and child tenants.

- **Provisioning**: provisioning activity based on the approver, datastore, requester, and virtual machine.

For a complete list of reportable fields in Red Hat CloudForms, see Appendix B, *Reportable Fields in Red Hat CloudForms*.

3.1. RUNNING REPORTS

There are two different ways to generate a report: by scheduling the report, or running the report manually by clicking the report generation button on the **Reports** page. Red Hat CloudForms uses interactive report generation so that reports are placed on a queue. A visual indicator of the reports status is shown. All reports are automatically saved so that they can be downloaded and analyzed later.

3.1.1. Generating a Single Report

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion. Navigate to the report you want to view by selecting the appropriate folder.

3. Click ⚙ (**Queue**).

4. The report generation is placed on the queue and its status shows in the reports page.

	Queued At	Run At	Source	User ID	Status
⏸	10/21/11 14:14:11 UTC		Requested by user	admin	Queued
✅	10/21/11 14:00:08 UTC	10/21/11 14:00:22 UTC	Requested by user	admin	Finished

5. Click ↻ (**Reload current display**) to update the status.

6. When a report has finished generating, click on its row to view it.

3.2. SCHEDULING A REPORT

You can view historical data by creating reports on a scheduled basis. In addition, scheduled reports can be emailed directly to users.

> **NOTE**
>
> You may need to disable, change the report filter, or change the frequency of a schedule. To do this, you will need to edit the schedule.

There are two ways to schedule a report:

Add a new Schedule through the **Reports** accordion:

1. Navigate to **Overview → Reports**.

2. Select a report from the **Reports** accordion and click ⚙ **Configuration**.

3. Click ⊕ **Add a New Schedule** or click the Schedules accordion and click ⚙ **Configuration** > ⊕ **Add a New Schedule**.

or

Add a new Schedule through the **Schedules** accordion:

1. Navigate to **Overview → Reports**.

2. Click the **Schedules** accordion.

3. Click ⚙ **Configuration**.

4. Click ⊕ **Add a New Schedule**.

Adding new Schedule details.

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

   - By default, **Active** is checked to enable the scan.

2. The **Report Selection** area is pre-populated if you added the schedule directly from the report. If you are adding from the schedule according, use the **Filter** drop downs to select the report that you want to schedule.

3. In the **Timer** area, click the **Run** drop down to specify how often you want the analysis to run. Your options after that will depend on which run option you choose.

   - Click **Once** to have the analysis run just one time.

   - Click **Daily** to run the analysis on a daily basis. You will be prompted to select how many days you want between each analysis.

   - Click **Hourly** to run the analysis hourly. You will be prompted to select how many hours you want between each analysis.

   - Type or select a date to begin the schedule in **Starting Date**.

   - Select a **Starting Time** based on a 24 hour clock in the Red Hat CloudForms appliance's Time Zone.

4. Use the **E-Mail after Running** feature to send an email after the report has been generated to selected users. The email will include a link to the report.

   - Check **Send an E-mail**

   - In **From (leave blank for default)** type in the sending email.

   - Use **Add a User**, to select a specific user. The user must have a valid email address entered under accounts.

   - Use **Add (enter manually)** to type in the address not registered to a User.

   - Then, click ⊕ (Add).

5. If you are sending an email after the report runs, then you can select further options under Email Options.

   - Check **Send if Report is Empty** if you want an email even if no records exist in the report.

   - Next to Attachments, check if you would like the report attached as a **Text**, **CSV**, or **PDF** file.

6. Click **Add**.

> **NOTE**
>
> See https://access.redhat.com/documentation/en/red-hat-cloudforms/4.7/paged/general-configuration in the *General Configuration* guide, to learn how to verify the address and validate outgoing email settings.

3.2.1. Modifying a Report Schedule

1. Navigate to **Overview → Reports**

2. Click the **Schedules** accordion and select the schedule you want to edit.

3. Click (**Configuration**), then click ![pencil] (**Edit this Schedule**).

4. Make the required changes.

5. Click **Save**.

3.2.2. Running a Scheduled Report Immediately

1. Navigate to **Overview → Reports**.

2. Click the **Schedules** accordion and select the schedule you want to run.

3. Click ![gear] (**Configuration**), then click ![gear] (**Queue up this Schedule to run now**).

3.3. VIEWING REPORTS

Once you have run a report, you can view it at any time after the first scheduled time has occurred.

1. Navigate to **Overview → Reports**.

2. Click the **Saved Reports** accordion or the **Reports** accordion.

3. Click on the instance of the report you want to view.

3.3.1. Changing Report Views

After running a report, you can customize your view in several ways.

Some reports can be viewed as charts as well as lists. This will depend on the type of data and on how the report has been created. Where applicable, you will see these additional buttons:

1. Navigate to **Overview → Reports**.

2. Click a completed report to view. Click one of the following buttons for the view you want:

   - Click ![graph icon] for **Graph View**.

   - Click ![hybrid icon] for **Hybrid View**.

   - Click ![tabular icon] for **Tabular View**.

3.3.1.1. Filtering Report Data

To filter the report information you see, use the ![filter icon] **Data View** button.

1. Navigate to **Overview → Reports**.

2. Click a completed report to view.

3. Click the ![filter icon] button for **Data View**, which opens the search field.

4. Select the criteria from the dropdown menu to filter by, and enter your search term.

5. Press **Enter** to search.

3.3.1.2. Sorting Reports by Column

In tabular and hybrid view, you can sort a list of reports by column.

1. Navigate to **Overview → Reports → Saved Reports**.

2. Click a report name to show a list of completed reports.

3. Click any blue column heading to sort the reports by the values in that column.

3.3.2. Downloading a Report

Download reports to analyze the data using other tools or to print the report.

When you click on one of the supplied reports, you are presented with a group of buttons to download the report in one of three formats or to view the report in a full screen.

1. Navigate to **Overview → Reports**.

2. Click the report you want to view.

3. Click on the row for the instance of the report you want to download. If the report needs to be generated, see Section 3.1, "Running Reports".

4. Click on the report download buttons for the type of export you want.

   - Click **(Download this report in text format)** to download as text.

   - Click **(Download this report in csv format)** to download as a comma-separated file.

   - Click **(Download this report in PDF format)** to download as PDF.

   - The report is automatically named with the type of report and date.

> **NOTE**
>
> Edit and delete buttons are only visible to administrators and super administrators. Edit and delete functions are only available to customer-created reports. The Red Hat CloudForms pre-configured reports cannot be edited or deleted, but they can be copied.

3.3.3. Showing a Report in Full Screen

View the report in full screen to examine report details. From full screen, you can also print the chart that accompanies a report.

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion and select the report that you want to view.
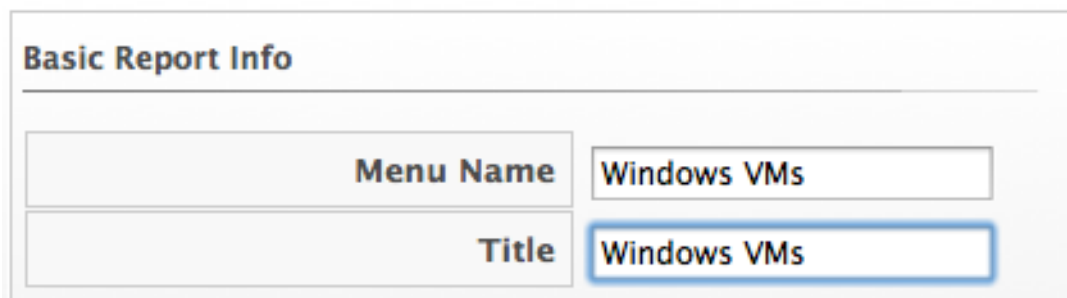
3. Click the **Saved Reports** tab and select a report.

4. On the **Saved Report** detail screen, click ![gear] (**Configuration**), and click ![document] (**Show full screen Report**).

The report will appear in full screen view. Click back on your browser to return to the reports explorer.

3.4. ADDING A REPORT

Add reports if the default reports do not include what you need or you want to narrow the scope of a report. For example, you may want a report that shows only Windows virtual machines.

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion.

3. Click ![gear] (**Configuration**), then click ![plus] (**Add a New Report**).

4. In the **Columns** tab, edit the **Basic Report Info** area.

- Type a unique name in **Menu Name** for how you want the report described in the menu list.

- Type the **Title** you want displayed on the report in title.

5. Add fields in the **Configure Report Columns** area.

Configure Report Columns

* Base the report on	VMs ⬍

Available Fields:

Active
Allocated Disk Storage
Annotation
Archived
Autostart
Busy
CPU – % Overallocated
CPU – Aggregate Usage Rate for Child VMs for Collected Intervals 30 Day Max (MHz)
CPU – Aggressive Recommendation Savings (%)
CPU – Aggressive Recommendation Savings

▼　▲

Selected Fields:

Name
Boot Time
OS Name

⤒ ▲ ▼ ⤓

* Caution: Changing these fields will clear all selected values below them!

- Use the **Base the report on** table dropdown to choose a table to get fields from.

> **NOTE**
>
> If you change the report base or the interval, all selections below will be reset.

- Select fields that you want in the report from **Available Fields**, and then click ⌄ **(Move selected fields down)**. In addition to the fields, you can also select any tags that you have created and assigned.

- Change the order of the fields in the report by clicking ⌃ **(Move selected fields up)** or ⌄ **(Move selected fields down)**.

6. Click on the **Consolidation** tab to consolidate and aggregate data points into maximum, minimum, average, and total. Specifically, this would be useful for analyzing performance data over a specific period of time.

> **NOTE**
>
> Please note that if you do this, you will not see individual records, but rather the calculation as a column header.

Group Records by up to 3 Columns

Column 1	Host : Name
Column 2	Asset Name
Column 3	<<< Nothing >>>

Note: Consolidating records will not show detail records in the report.

- Select the columns to group by.

Specify Calculations of Numeric Values for Grouped Records

Column Name	Calculations	Selected
CPU – Usage Rate for Collected Intervals (%)	Check Options	
CPU – Usage Rate for Collected Intervals (MHz)	Check Options	
CPU – Total – from VM Analysis (MHz)	Check Options	

- For each numeric field selected in the report, you can click the dropdown under **Calculations**. Check the calculations you want to use.

7. Click on the **Formatting** tab to set the size of paper for a PDF and column header format.

- From the **PDF Output** area, select the page size from the dropdown.

PDF Output

Page Size	US Letter – 8.5in x 11.0in

A0 – 841mm x 1189mm
A1 – 594mm x 841mm
A2 – 420mm x 594mm
A3 – 297mm x 420mm
A4 – 210mm x 297mm (default)
US Executive – 7.25in x 10.5in
US Folio – 8.5in x 13.0in
US Government – 8.0in x 11.0in
US Ledger – 17.0in x 11.0in
US Legal – 8.5in x 14.0in
US Letter – 8.5in x 11.0in
US Statement – 5.5in x 8.5in
US Tabloid – 11.0in x 17.0in

- From **Specify Column Headers and Formats**, type the text you want displayed for each field. For each numeric field, you can also set the numeric format.

Specify Column Headers and Formats

Column Name	Header	Format
Host : Name	VM's Host Name	<None>
Activity Sample – Timestamp (Day/Time)	Activity Sample – Timestamp (Day	Date/Time (M/D/Y H:M:S Z)
Asset Name	Asset Name	<None>
CPU – Usage Rate for Collected Intervals (%)	CPU – Usage Rate for Collected Int	Percent, 1 Decimal (99.0%)
CPU – Usage Rate for Collected Intervals (MHz)	CPU – Usage Rate for Collected Int	Megahertz Avg (12.1 Mhz)
CPU – Total – from VM Analysis (MHz)	CPU – Total Available – from VM A	Megahertz (12 Mhz)

8. Click on the **Styling** tab to change the color of the text or the background for a row based on a condition.

Specify Column Styles

Column Name	Style	If	
Host : Name	\<Normal>	Default	
Activity Sample – Timestamp (Day/Time)	\<Normal>	Default	
Asset Name	\<Normal>	Default	
CPU – Usage Rate for Collected Intervals (%)	Blue Text	>= 80 %	**Sample**
	\<Normal>	Default	
CPU – Usage Rate for Collected Intervals (MHz)	\<Normal>	Default	
CPU – Total – from VM Analysis (MHz)	\<Normal>	Default	

* Style "If" conditions are evaluated top to bottom for each column

- Use **Style** to select the format for the value. You can choose to change the text color or the background.

- Use **If** to create a conditional statement for the style.

9. Click on the **Filter** tab to set filters for the data displayed in the report. There are two types of filters: the first is the **Record Filter** which is the primary filter of the main tables records, the second is also a **Display Filter**, which is a secondary filter of rows based on the fields of the child table. Click ✏ in the appropriate area to use the expression editor to choose the appropriate options for your criteria. Based on what you choose, different options will show.

- Click **Field** to create criteria based on field values.

- Click **Count of** to create criteria based on the count of something, such as the number of snapshots for a virtual machine, or the number of virtual machines on a host.
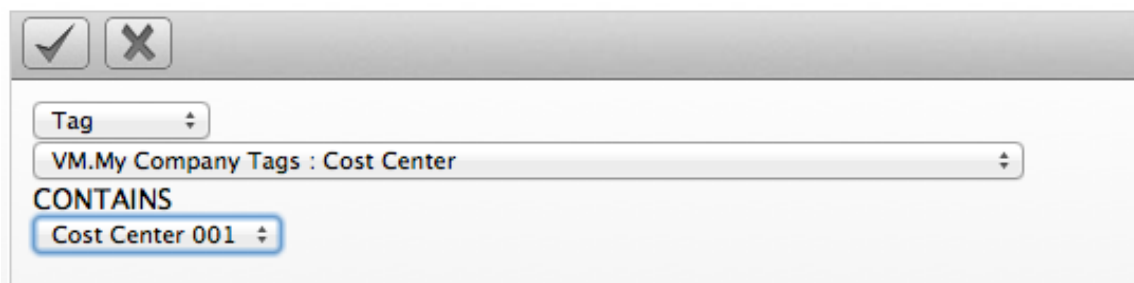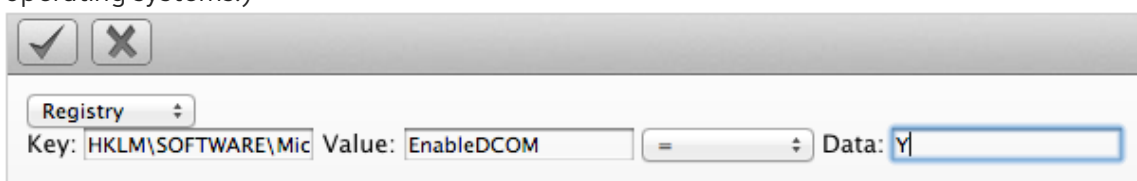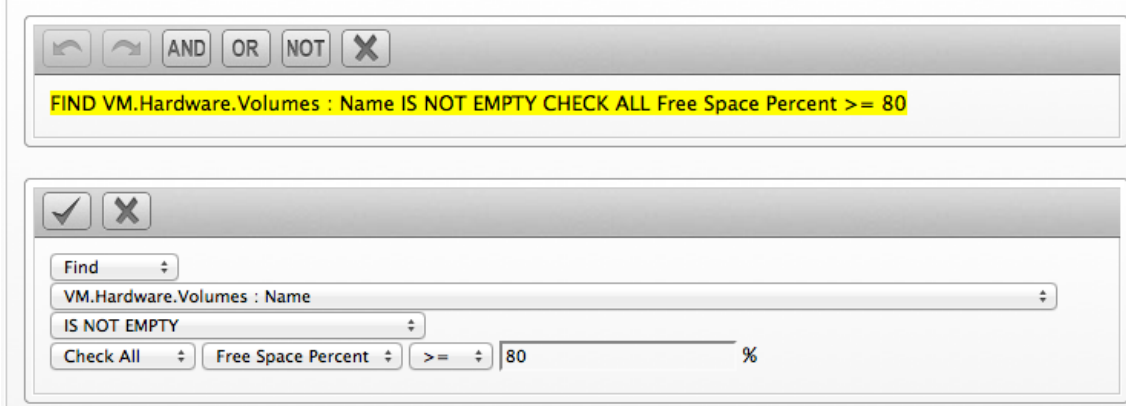
- Click **Tag** to create criteria based on tags assigned to your resources. For example, you may want to check the power state of a virtual machine or see if it is tagged as production.

- Click **Registry** to create criteria based on registry values. For example, you may want to check if DCOM is enabled on a Windows system. (Note that this applies only to Windows operating systems.)

- Click **Find** to seek a particular value, and then check a property. For example, finding the **Admin** account and checking that it is enabled.

- Click ✔ **(Commit Expression Element Changes)** to add the expression.

> **NOTE**
>
> The filters that you apply will show at the bottom of the report so that you know which filters have been applied.

10. Click on the **Summary** tab to select sort order, sort type, groupings, and group calculations for the report. **Summary** groups rows of data based on the sort breaks. You can only sort by fields that display in the report.

- Set the primary sort in **Sort the Report by**.

- Set the next sorts in **Within Above Field, Sort By**.

- Select the type of sort, ascending or descending, in **Sort order**.

- In **Show Sort breaks**, select **Yes** to show the sort breaks, **Counts** to show sort breaks with the count, or **No** for no sort breaks.

- For any numeric field, you can select to show minimum, average, maximum, and total in the sort break.

11. Click on the **Charts** tab to create a chart for the report. This is not required.

- Use **Choose a chart type** to select a type of chart. Note that some charts may not produce the result you are looking for based on the types of fields in the report and its sort order.

- If you only want to see the top values, select the number of top values from **Top values to show**.

- If you want to see the total number of values that are not categorized or evaluated against others, check **Sum 'Other' values**.

12. Click on the **Timeline** tab to select a timeline for the report. You must have a field of time or date format to use this feature.

Timeline Settings

Base Timeline on	Activity Sample – Timestamp (Day/Time)
First band unit:	Hour
Second band unit:	Day
Third band unit:	<<< Nothing >>>

Event to position at:	Last
Show events from last:	7 Days

- Use **Base Timeline on** to select a column in date or time format for the report.

- Select a unit of time for the first band in **First band unit**.

- Select a unit of time for the second band in **Second band unit**.

- Select a unit of time for the third band in **Third band unit**.

- Select an **Event to position at**.

- Select the range for the event to position from **Show events from last**.

13. Click the **Previews** tab to see a sample of your report.

14. When you have the report that you want, click **Add** to create the new report.

> **NOTE**
>
> After the new report is created, to make the report accessible from the report menu, you must add it to a report menu.

3.5. COPYING A REPORT

Use this feature to copy a report that is similar to one that you want to create. By copying a report, you only need to make minor edits instead of creating an entirely new report.

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion and select the report that you want to copy.

3. Click (**Configuration**), then click (**Copy this report**).

4. On the Columns tab, edit the **Basic Report Info** area to include a new **Menu Name**. Each **Menu Name** must be unique.

5. Make any other changes you need. See "Adding a Report" for details on the changes you can make.

6. Click **Add**.

3.6. EDITING A REPORT

If you find that a report is not giving you the data that you need, you can edit it after it has been created. Note that only reports that you have created can be modified. Only administrators and super administrators of Red Hat CloudForms can add, copy, edit, and delete reports.

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion and select the report you want to edit.

3. Click ⚙ (**Configuration**), ✏ (**Edit this Report**).

4. Make any changes you need.

5. Click **Save**.

3.7. DELETING A REPORT

Delete reports when you find that they are no longer useful. Only administrators and super administrators of Red Hat CloudForms can add, copy, edit, and delete reports.

> **NOTE**
>
> Please note that only customer-created reports can be deleted.

1. Navigate to **Overview → Reports**.

2. Click the **Reports → All Reports** accordion and select the report you want to delete.

3. Click ⚙ (**Configuration**), 🗑 (**Delete this Report from the Database**).

> **NOTE**
>
> The **Delete this report from the Database** option will only appear on reports you have created. Default reports cannot be deleted.

4. Click **OK**.

3.8. IMPORTING A REPORT

Reports are stored in the VMDB; however, you can share customized reports among VMDBs. To do this use the export and import feature of reports.

1. Navigate to **Overview → Reports**.

2. Select **Custom Reports** from the **Import/Export** accordion.

3. In the **Import** area, click **Browse** to select an import file.

4. To overwrite an existing report with the same menu name, select **Overwrite existing reports**.

5. Click **Upload** to import the report to the **Company-Custom** folder.

3.9. EXPORTING A REPORT

Reports are stored in the VMDB. Custom reports may be shared among VMDBs. Share your custom reports using the export and import features.

1. Navigate to **Overview → Reports**.

2. Select **Custom Reports** from the **Import/Export** accordion.

3. In the **Export** area, select the reports to export.

4. Click **Export**.

3.10. MANAGING REPORT MENUS

By default, all account roles have the same reports available. If you have the super administrator or administrator role, you can customize the accordions, the folders within them, and the locations of your reports.

> **NOTE**
>
> When you create your own report, the report is not automatically available. You must add it to a report menu.

The report menu is a hierarchical structure that consists of the following components:

- The top level, under which the accordions are shown. (Top level is only displayed when you are in the report menu editor.)

- Accordions that are general categories for the reports. The defaults supplied are **Configuration Management**, **Migration Readiness**, **Operations**, **VM Sprawl**, **Relationships**, **Events**, **Performance by Asset Type**, **Trending**, **Tenants**, **Provisioning**, and **Running Processes**.

- Folders that are used to further organize reports within an accordion. For example, under the **Configuration Management**, there are folders for virtual machines, hosts, and other virtual infrastructure components.

- Reports that are stored directly in the folders.

3.10.1. Managing Report Menu Accordions

1. Navigate to **Overview → Reports**.

2. Click the **Edit Report Menus** folder.

3. Click the role whose menus you want to customize.

4. Click on **Top Level** to organize, add, and delete folders.

- Click ⏶ **(Move selected folder to top)** to move the folder to the top of the list.

- Click ⌃ **(Move selected folder up)** to move the folder up.

- Click ⌄ **(Move selected folder down)** to move the folder down.

- Click ⌄⌄ **(Move selected folder to bottom)** to move the folder to the bottom of the list.

- Click ▬ **(Delete selected folder and its contents)** to delete a folder.

- Click ✚ **(Add folder to selected folder)** to add an folder.

> **NOTE**
>
> If you are creating a new folder, **Top Level** must be selected under **Reports**. Be sure to select the folder you want to create a subfolder for on the left pane. To name the folder, double-click on **New Folder**, then click on **Top Level** in the **Reports** area.

5. When you are finished adding folders, click **Commit**. To revert, click **Discard**.

6. Click **Save**.

The new folder is added, and you can add folders in which to store reports to it. You can also organize the reports into folder that are meaningful to you.

3.10.2. Managing Report Menu Folders

1. Log in to the console as a user who is assigned either the super administrator or administrator account role.

2. Navigate to **Overview → Reports**.

3. Click the **Edit Report Menus** accordion.

4. Click the role whose menus you want to customize.

5. Click on the accordion name you want to organize or add folders to.

- Click ⌃⌃ **(Move selected folder to top)** to move the folder to the top of the list.

- Click ⌃ **(Move selected folder up)** to move the folder up.

- Click ⌄ **(Move selected folder down)** to move the folder down.

- Click ⮟ (Move selected folder to bottom) to move the folder to the bottom of the list.

- Click ▬ (Delete selected folder and its contents) to delete an folder.

- Click ✚ (Add subfolder to selected folder) to add a folder. When creating a subfolder, be sure to select the folder that you want the subfolder to show under. To name the folder, double-click on **New Folder**.

6. When you are finished making changes click **Commit**. To revert, click **Discard**.

7. Click **Save**.

> **NOTE**
>
> Only reports that are not already in another folder can be assigned.

3.10.3. Organizing Reports in Report Menus

1. Log in to the console as a user who is assigned either the Super Administrator or Administrator Account Role.

2. Navigate to **Overview → Reports**.

3. Click the **Edit Report Menus** accordion.

4. Click the role whose menus you want to customize.

5. Expand the **Report** accordion and menus using the triangles to the left of the item name until you are able to select the subfolder where you want to put reports.

6. Choose one of the following actions:

   - To add a report, select a report from the **Available Reports** area on the right and click

     ‹ (Move selected reports left).

   - To remove a report from a folder, select the report from the **Selected Reports** area and click › (Move selected reports right).

   - To move a report to the top of the folder, select the report and click ⮝ (Move selected reports to top).

   - To move a report up one place in the folder, select the report and click ⌃ (Move selected reports up).

- To move a report down one place in the folder, select the report and click ⌄ (**Move selected reports down**).

- To move a report to the bottom of the folder, select the report and click ⌄⌄ (**Move selected reports to bottom**).

7. When you are finished making changes click **Commit**.

8. Click **Save**.

The changes are made. The next time a user with this Account Role logs in, the new report and menu structure will be displayed.

CHAPTER 4. CHARGEBACK

The chargeback feature allows you to calculate monetary virtual machine charges based on owner or company tag. To use this feature you must be collecting capacity and utilization data.

For information on server control settings and capacity & utilization collection settings, see *General Configuration*.

4.1. CHARGEBACK RATES

Red Hat CloudForms provides a default set of rates for calculating chargeback costs, but you can create your own set of computing and storage costs by navigating to **Overview → Chargeback** and clicking the **Rates** accordion.

You can configure chargeback rates for various resources by selecting either **Compute** or **Storage** in the **Rates** accordion. **Compute** sets chargeback rates for CPU, disk I/O, fixed compute cost, memory, and network I/O items, while **Storage** sets chargeback rates for fixed storage cost and disk storage.

Chargeback costs are computed using a set formula based on hourly cost per unit and hourly usage.

Chargeback can be calculated in the following currencies:

- Australian dollar (AUD)

- Brazilian real (BRL)

- Swiss franc (CHF)

- Chinese yuan renminbi (CNY)

- Euro (EUR)

- Pound sterling (GBP)

- Indian rupee (INR)

- Japanese yen (JPY)

- South Korean won (KRW)

- Mexican peso (MXN)

- Norwegian krone (NOK)

- Russian ruble (RUB)

- Swedish krona (SEK)

- Turkish lira (TRY)

- United States dollar (USD)

- South African rand (ZAR)

Chargeback rates can be assigned at a single rate or by tiers, where rates are assigned in ranges depending on level of usage.

Additionally, chargeback can be calculated at one fixed rate, or by a combination of fixed and variable rates per tier. Fixed rates are charged once per unit of time, and the variable rate is calculated by the level of usage multiplied by the number of resources used in a unit of time.

4.1.1. Memory Used Cost

Calculating the Memory Used Cost in dollars ($) for a day can be expressed in the following ways:

- Memory allocation per hour (in MB) * Hourly Allocation cost per megabyte * Number of Memory Allocation metrics available for the day

- Sum of Memory allocation for the day (in MB) * Hourly Allocation cost per megabyte

- Sum of Memory allocation for the day (in MB) * Daily Allocation cost per megabyte / 24

Memory costs can be measured in B, KB, MB, GB, or TB.

> **Example 4.1. Memory Used Cost**
>
> In a scenario where 9.29 GB of memory is used in a day with the chargeback rate set at one dollar ($1) per megabyte per day, the Memory Used Cost would be $396.42.
>
> - 9.29 GB = 9514.08 MB
>
> - 9514.08 MB * $1 (per MB per day) = $9514.08
>
> - $9514.08 / 24 = $396.42 Memory Used Cost

4.1.2. CPU Total Cost

The CPU Total Cost is defined as the number of virtual CPUs over the selected interval (hour, day, week, month). CPU costs can be measured in units of Hz, KHz, MHz, GHz, or THz, as specified when creating a chargeback rate.

> **Example 4.2. CPU Total Cost**
>
> In a scenario where 16 CPUs are used in a day with the chargeback rate set at one dollar per CPU per day, the CPU Total Cost would be $16.
>
> - 16 CPUs * $1 (per CPU per day) = $16 CPU Total Cost

4.1.3. CPU Used Cost

The CPU Used Cost is defined as the average CPU used in MHz over the selected rate interval (hour, day, week, month). CPU Used Cost is not supported for containers providers.

> **Example 4.3. CPU Used Cost**
>
> In a scenario where 2.5 GHz is used in a day with the chargeback rate set at $0.01 per MHz per day, the CPU Used Cost would be $25.
>
> - 2.5 GHz = 2500 MHz

- 2500 MHz * $0.01 (per MHz per day) = $25 CPU Used Cost

4.1.4. Storage Allocated Cost

The Storage Allocated Cost is defined as the Allocated Disk Storage in bytes over the selected rate interval (hour, day, week, month). Storage costs can be measured in B, KB, MB, GB, or TB.

Example 4.4. Storage Allocated Cost

In a scenario where 500 GB are used in a day with the chargeback rate set at $0.10 per GB per day, the Storage Allocated Cost would be $50.

- 536,870,912,000 bytes = 500 GB

- 500 GB * $0.10 (per GB per day) = $50 Storage Allocated Cost

4.1.5. Storage Total Cost

The Storage Total Cost is defined as the Used Disk Storage in bytes over the selected rate interval (hour, day, week, month).

Example 4.5. Storage Total Cost

In a scenario where 250 GB are used in a day with the chargeback rate set at $0.10 per GB per day, the Storage Total Cost would be $25.

- 268,435,456,000 bytes = 250 GB

- 250 GB * $0.10 (per GB per day) = $25 Storage Total Cost

4.1.6. Storage Used Cost

The Storage Used Cost is defined as the Used Disk Storage in bytes over the selected rate interval (hour, day, week, month).

Example 4.6. Storage Used Cost

In a scenario where 250 GB are used in a day with the chargeback rate set at $0.10 per GB per day, the Storage Used Cost would be $25.

- 268,435,456,000 bytes = 250 GB

- 250 GB * $0.10 (per GB per day) = $25 Storage Used Cost

NOTE

The following chargeback rates are not supported for containers providers:

- Allocated CPU count

- Used CPU

- Used disk I/O

- Allocated memory

4.2. CREATING CHARGEBACK RATES

Red Hat CloudForms allows you to create your own set of computing and storage costs to use for billing.

Chargeback rates can be configured at a single rate or in tiers, where one rate is assigned to one usage range, and another rate is assigned to a different usage range. You can also assign fixed and variable rates per tier if desired.

There are three methods to create a chargeback rate: editing or copying an existing rate, or creating a new rate.

4.2.1. Editing a Chargeback Rate

1. Navigate to **Overview → Chargeback**.

2. Click the **Rates** accordion and select an existing rate from under **Compute** or **Storage**.

3. Click ⚙ (**Configuration**) and 🖊 (**Edit this Chargeback Rate**).

4. Edit the rate as desired.

5. Click **Save** at the bottom of the screen to save the chargeback rate.

4.2.2. Copying a Chargeback Rate

1. Navigate to **Overview → Chargeback**.

2. Click the **Rates** accordion and select an existing rate from under **Compute** or **Storage**.

3. Click ⚙ (**Configuration**) and 🗐 (**Copy this Chargeback Rate**).

4. Edit the rate as desired.

5. Click **Add** at the bottom of the screen to save the chargeback rate.

4.2.3. Adding a Chargeback Rate

The following example shows a scenario to create a new chargeback rate for allocated CPU usage in Euro per week measured by the following tiers, calculating both fixed and variable costs:

Example: Creating a tiered chargeback rate

- Tier 1: 0-2 CPU = 1 Euro fixed + 0.5 Euro variable

- Tier 2: 2-4 CPU = 0.8 Euro fixed + 0.4 Euro variable

- Tier 3: 4 CPU and above = 0.6 Euro fixed + 0.2 Euro variable

To configure this rate:

1. Navigate to **Overview → Chargeback**.

2. Click the **Rates** accordion and select **Compute** to create a CPU chargeback rate.

   a. Click ⚙ (**Configuration**) and ⊕ (**Add a new Chargeback Rate**).

3. Type in a **Description** for the chargeback rate.

4. Select **Euro** for currency.

5. In **Rate Details**, select the time, and fill in the range and rates for each tier. Alternatively, a single chargeback rate can be configured by using only one tier. For this example, in **Allocated CPU Count**, configure:

   a. For tier 1, set:

      i. In **Per Time**, select **Hourly**.

      ii. In **Range**, specify the CPU range: **Start** remains as **0.0** and **Finish** is **2**.

      iii. In **Rate**, specify **1** for the **Fixed** rate, and **0.5** for the **Variable** rate.

      iv. Click **Add** under **Actions** to create a new, blank **Allocated CPU Count** line for tier 2.

   b. For tier 2, set:

      i. In the new line in **Range**, specify the CPU range: **Start** is **2** and **Finish** is **4**.

      ii. In **Rate**, specify **0.8** for the **Fixed** rate, and **0.4** for the **Variable** rate.

      iii. Click **Add** under **Actions** to create another new, blank **Allocated CPU Count** line for tier 3.

   c. For tier 3, set:

      i. In the new line in **Range**, specify the CPU range: **Start** is **4** and leave **Finish** as **Infinity**.

      ii. In **Rate**, specify **0.6** for the **Fixed** rate, and **0.2** for **Variable**.

Compute Chargeback Rates

Basic Info

Description: CPU Count - Paris

Currencies

Select currency: € [Euro]

Rate Details

*Caution: The value Range end will not be included in the tier.

Group	Description	Per Time	Per Unit	Range		Rate		Actions	Currency
				Start	Finish	Fixed	Variable		
Cpu	Allocated CPU Count	Hourly		0.0	2	1.0	0.5	Add	EUR
				2	4	0.8	0.4	Delete	
				4	Infinity	0.6	0.2	Delete	
CPU	Used CPU	Hourly	MHz	0.0	Infinity	0.0	0.02	Add	EUR
Cpu Cores	Used CPU Cores	Hourly		0.0	Infinity	1.0	0.02	Add	EUR
Disk I/O	Used Disk I/O	Hourly	KBps	0.0	Infinity	0.0	0.005	Add	EUR
Fixed	Fixed Compute Cost 1	Hourly		0.0	Infinity	0.0	0.0	Add	EUR
Fixed	Fixed Compute Cost 2	Hourly		0.0	Infinity	0.0	0.0	Add	EUR
Memory	Allocated Memory	Hourly	MB	0.0	Infinity	0.0	0.0	Add	EUR
Memory	Used Memory	Hourly	MB	0.0	Infinity	0.0	0.02	Add	EUR
Network I/O	Used Network I/O	Hourly	KBps	0.0	100.0	0.5	0.0	Add	EUR
				100.0	Infinity	0.5	0.005	Delete	

6. Click **Add** at the bottom of the screen to save the chargeback rate.

The rate can now be assigned.

4.3. ASSIGNING CHARGEBACK RATES

Red Hat CloudForms allows you to assign chargeback rates by choosing from **Compute** and **Storage**. You can assign either a custom chargeback rate, or a default rate for Compute, Storage, or Container Images.

> **NOTE**
>
> The **Default Container Image Rate** can be useful when applied to a containers provider with multiple images, where only some of the images have a chargeback rate assigned; the images without a custom rate will then use the default rate.

1. Navigate to **Overview → Chargeback**.

2. Click the **Assignments** accordion, and click either **Compute** or **Storage**.

4.3.1. Configuring Chargeback for Compute

Use **Compute** to assign a compute chargeback rate. You can assign chargeback rates to **The Enterprise**, **Selected Cloud/Infrastructure Providers**, **Selected Clusters/Deployment Roles**, **Selected Containers Providers**, **Tagged VMs and Instances**, **Tagged Images**, **Labeled Images**, or **Tenants**.

1. From the **Basic Info** area, use the **Assign To** list to select a type of assignee to assign the rate set to.

2. For **Tagged VMs and Instances** and **Tagged Images** select a **Tag Category**.

3. For each item to set, select the chargeback rate to use. The options displayed vary based on the type you select.

4. Click **Save**.

4.3.2. Configuring Chargeback for Storage

Use **Storage** to assign a storage chargeback rate. You can assign chargeback rates to **The Enterprise**, **Selected Datastores**, **Tagged Datastores**, or **Tenants**.

1. From the **Basic Info** area, use the **Assign To** list to select a type of assignee to assign the rate set to.

2. For **Tagged Datastores** select a **Tag Category**.

3. For each item to set, select the chargeback rate to use. The options displayed vary based on the type you select.

4. Click **Save**.

The rate is assigned. The next time you generate a chargeback report, these values will be used.

> **NOTE**
>
> When viewing chargeback, there is a rate for a virtual machine for the number of the CPUs. The chargeback for this parameter is calculated based on when the virtual machine is running. If the virtual machine is not running, then it is not charged for CPU allocation.

4.4. CREATING A CHARGEBACK REPORT

Red Hat CloudForms allows you to create chargeback reports to monitor costs you charged.

> **NOTE**
>
> When creating a chargeback report, the **Total Cost** is the complete cost for the chosen workload over the selected time frame. Total costs are calculated by adding the calculated costs for the amount used and the amount allocated for the time frame. When adding cost related columns to display on your report, be sure to select all rate assigned fields to present a comprehensive accounting of total costs.

1. Navigate to **Overview → Reports**.

2. Click the **Reports** accordion.

3. Click ⚙ (Configuration), ⊕ (Add a new Report).

4. On the **Columns** tab, fill out the **Basic Report Info** area.

   - Type a unique name in **Menu Name** for how you want the report described in the menu list.

   - Type the **Title** to display on the report.

5. Add fields in the **Configure Report Columns** area.

   - From the **Base the report on** list, select **Chargebacks**.

- Select the fields to include in the report from the **Available Fields** list, then click ⌄ **(Move selected fields down)**. In addition to the fields, you can also select any tags that you have created and assigned.

- Change the order of the fields in the report by clicking ⌃ **(Move selected fields up)** or

  ⌄ **(Move selected fields down)**].

6. Click the **Formatting** tab to set the size of paper for a PDF and column header format.

   - From the **PDF Output** area, select the page size from the **Page Size** list.

   - From **Specify Column Headers and Formats**, type the text to display for each field. For each numeric field, you can also set the numeric format.

7. Click the **Filter** tab to set filters for the data displayed in the report.

   - From **Chargeback Filters**, select how you want the costs to show, the tag category, the tag, and how you want the items grouped.

   - From **Chargeback Interval**, select the time interval. You must have a full interval worth of data in order to select an option other than **Partial** in the **Daily Ending With** list.

8. Click the **Preview** tab, and then **Load** to see what the report will look like.

9. When you are satisfied that you have the report that you want, click **Add** to create the new report.

The new report is created. To make the report accessible from the **Report** menu, you must add it to a report menu.

CHAPTER 5. OPTIMIZATION

Optimization providers administrators with a series of reports to provide meaningful information on areas to optimizing resources such as CPU, memory and disk and identify potential savings.

5.1. RUNNING AN OPTIMIZATION REPORT

1. Navigate to **Overview → Optimization**.

2. Click **Queue Report** for each report type you choose to run.

## 5.2. VIEWING OPTIMIZATION REPORTS

1. Navigate to **Overview → Optimization**.

2. Click on the report type you wish to view details for.

3. On the report type screen, click on a report instance to view details.

4. Sort reports by clicking on the field names at the top of the report.

CHAPTER 6. ALERTS

6.1. ASSIGNING THE NOTIFIER ROLE

1. Click (**Configuration**).

2. Click the **Settings** accordion, and select the Red Hat CloudForms server.

3. From the **Server Control** tab, select the **Notifier** role.

   Server Roles
   - ☐ Automation Engine
   - ☐ Capacity & Utilization Coordinator
   - ☐ Capacity & Utilization Data Collector
   - ☐ Capacity & Utilization Data Processor
   - ☑ Database Operations
   - ☐ Database Synchronization
   - ☑ Event Monitor
   - ☑ Notifier
   - ☑ Provider Inventory
   - ☑ Provider Operations
   - ☐ RHN Mirror
   - ☑ Reporting
   - ☑ Scheduler
   - ☐ SmartProxy
   - ☑ SmartState Analysis
   - ☑ User Interface
   - ☑ Web Services
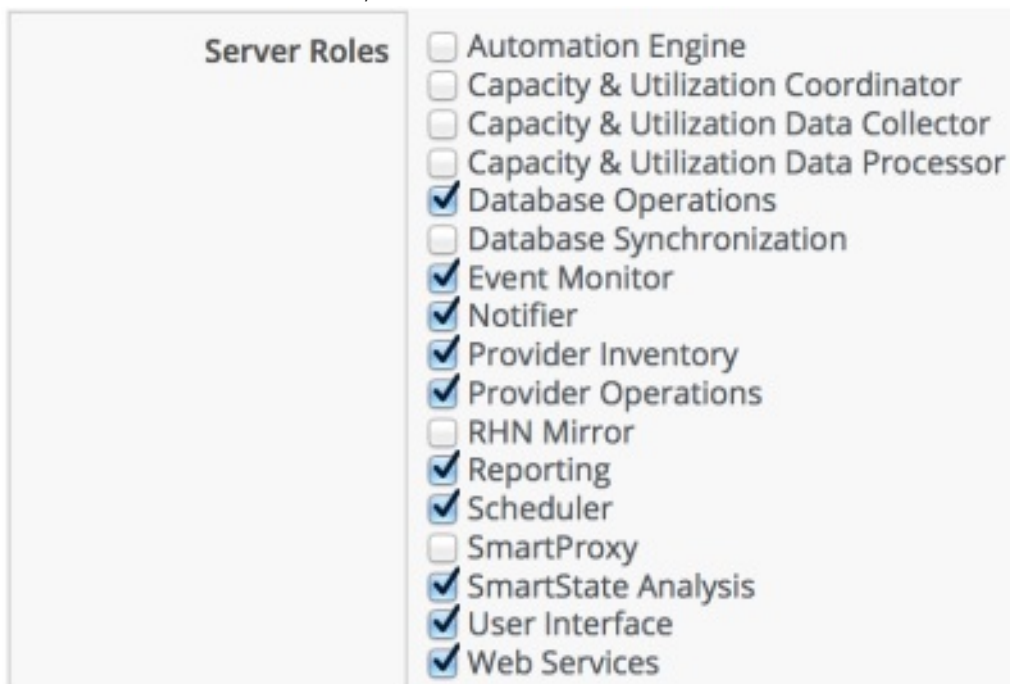
4. Click **Save**.

6.2. CREATING AN ALERT

This section outlines the basic procedure to create an alert.
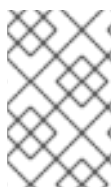
> **NOTE**
>
> To send emails or SNMP traps from the Red Hat CloudForms server, you must enable the **Notifier** server role and set up SMTP email or SNMP traps. For more information, see *General Configuration*.

1. Navigate to **Control → Explorer**.

2. Click the **Alerts** accordion, then click ![gear] (**Configuration**), ![plus] (**Add a New Alert**).

3. Enter the basic details of the alert:

   a. Enter a description in the **Description** field.

   b. Select the **Active** check box to enable the alert after creation.

   c. Select the severity level from the **Severity** list.

d. Select the inventory item on which to base the alert from the **Based On** list.

e. Select the type of event that triggers the alert from the **What to Evaluate** list.

f. Select the frequency with which to be notified if the event log threshold is reached from the **Notification Frequency** list.

4. Configure the parameters of the alert.

> **NOTE**
>
> The available parameters depend on the options you selected in the **Based On** and **What to Evaluate** lists. See later sections for the details of these parameters.

5. Optionally, select **Send an E-mail** to configure options so that an email is sent when the alert is triggered:

a. Enter the email address from which to send the email in the **From** field.

b. Select a user from the **Add a User** list to add a user registered in Red Hat CloudForms. The user must have a valid email address entered under accounts.

c. Enter the email address of a user in the **Add (enter manually)** field and click ✚ to add a user not registered in Red Hat CloudForms.

6. Optionally, select **Send an SNMP Trap** to configure options so that an SNMP trap is sent when the alert is triggered:

a. Enter the IP addresses of the hosts to send the trap in the **Host** fields.

b. Select the version of SNMP to use from the **Version** list:

i. If you select **v1**, enter a trap number in the **Trap Number** field. Enter **1**, **2**, or **3**, based on the appropriate suffix number in Table 6.1, "SNMP Trap Identifiers".

ii. If you select **v2**, enter a trap object ID in the **Trap Object ID** field. Enter **info**, **warning**, or **critical** based on the values in Table 6.1, "SNMP Trap Identifiers".

Table 6.1. SNMP Trap Identifiers

| Object ID	Suffix Number Added to PEN	PEN with the Suffix Added
info	1	1.3.6.1.4.1.33482.1
warn, warning	2	1.3.6.1.4.1.33482.2
crit, critical, error	3	1.3.6.1.4.1.33482.3

7. Optionally, select **Send a Management Event** to trigger an automation event:

a. Enter the name of the event that exists in the **Process/Event Class** in the **Event Name** field.

8. Click **Add**.

6.3. CREATING A HARDWARE RECONFIGURED ALERT

Use a hardware reconfigure alert to detect changes to the amount of memory or the number of CPUs on a virtual machine.

1. Navigate to **Control → Explorer**.

2. Click the **Alerts** accordion, then click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

3. In the **Info** area:

   - Type in a description for the alert.

   - From **Based On**, select **VM and Instance**.

   - From **What to Evaluate**, select **Hardware Reconfigured**.

   - In **Notification Frequency**, select how often you want to be notified if hardware reconfiguration is detected.

4. From **Hardware Attribute**, select Number of CPUs. From the next dropdown, select **Decreased**.

5. After setting the parameters, select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

6. Click **Add**.

6.4. CREATING A NORMAL OPERATING RANGE ALERT

Normal operating range alerts enables you to be notified when the normal operating range is exceeded, or falls below for a period of time from 1 minute to 2 hours. Capacity and utilization must be enabled for normal operating ranges to be calculated. See *General Configuration* for more information.

1. Navigate to **Control → Explorer**.

2. Click the **Alerts** accordion, then click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

3. In the **Info** area:

   - Type in a **Description** for the alert.

   - From **Based On**, select **VM and Instance**.

   - For **What to Evaluate**, select **Normal Operating Range**.

   - In **Notification Frequency**, select how often you want to be notified if the performance threshold is reached.

4. Set the threshold in the **Normal Operating Range Parameters** area.

- From **Performance Field**, select the field to check and whether you want to be notified if the field is exceeded or fell below.

- In **Field Meets Criteria for**, select the amount of time that the threshold requires to be met to trigger the alert.

5. After setting the parameters, you then select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process. See Section 6.2, "Creating an Alert".

6. Click **Add**.

6.5. CREATING A REAL TIME PERFORMANCE ALERT

Real Time Performance alerts enables you to be notified immediately when a performance threshold has been met for a virtual machine, host, or cluster. Capacity and Utilization must be enabled for performance thresholds to be detected. See *General Configuration* for more information.

1. Navigate to **Control → Explorer**.

2. Click the **Alert** accordion, then click (**Configuration**), (**Add a new Alert**).

3. In the **Info** area:

- Type in a **Description** for the alert.

- From **Based On**, select **VM and Instance**.

- For **What to Evaluate**, select **Real Time Performance**.

- In **Notification Frequency**, select how often you want to be notified if the performance threshold is reached.

4. Set the threshold in the **Real Time Performance Parameters** area.


- From **Performance Field**, select the field to check and any other parameters required for that field.

- In **And is Trending**, select **Don't Care** if it does not matter how the performance metric is trending. Otherwise, choose from the possible trending options.

- In **Field Meets Criteria for**, select the amount of time that the threshold requires to be met to trigger the alert.

- Set **Debug Tracing** to true only when directed to do so by Red Hat Support. This provides an extremely detailed level of logging and can result in many more log lines being written.

5. After setting the parameters, you then select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

6. Click **Add**.

6.6. CREATING AN HOURLY PERFORMANCE ALERT

Hourly performance alerts enable you to be notified immediately when an hourly performance threshold has been met for a cluster. Capacity and Utilization must be enabled for performance thresholds to be detected. See *General Configuration* for instructions.

1. Navigate to **Control → Explorer**.

2. Click the **Alerts** accordion.

3. Click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

4. In the **Info** area:

- Type in a **Description** for the alert.

- From **Based On**, select **Cluster**.

- For **What to Evaluate**, select **Hourly Performance**.

- In **Notification Frequency**, select how often you want to be notified if threshold is met.

5. In the **Hourly Performance Parameters** area select performance field and the criteria. You can also select options from the **And is Trending** dropdown box and whether the **Debug Tracing** is true or false.

6. After setting the parameters, you then select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

7. Click **Add**.

6.7. CREATING A HOSTD LOG THRESHOLD ALERT

Use the **hostd Log Threshold** to send a notification when certain items are found in the event logs for a host. A default analysis profile with event log items is required for this feature. The following example shows steps to check the host's log for a failure to validate a virtual machine's IP address.

1. Navigate to **Control → Explorer**.

2. Click the **Alert** accordion.

3. Click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

4. In the **Info** area:

   Info

 | | |
   |---|---|
   | Description | |
   | Active | ☑ |
   | Based On | Host / Node ⌄ |
   | What to Evaluate | Hostd Log Threshold ⌄ |
   | Notification Frequency | 10 Minutes ⌄ |

   - Type in a **Description** for the alert.

   - From **Based On**, select **Host/Node**.

   - For **What to Evaluate**, select **Hostd Log Threshold**.

   - In **Notification Frequency**, select how often you want to be notified if the log item is detected.

5. In the **Hostd Log Threshold Parameters** area, select the parameters for the event log message. You can set a threshold for a filter, level, or message source.

   - Use **Message Filter** to look for specific text in a message. Use **Message Level** to filter based on message level. Red Hat CloudForms reports on the specified level and above. Use **Message Source** to filter log messages based on its source.

   - Set **How Far Back to Check** in days you want to look for this message.

   - If you only want an alert triggered when the log message has occurred a certain number of times, type the number in **Event Count Threshold**.

6. After setting the parameters, select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

7. Click **Add**.

6.8. CREATING A VMWARE ALARM ALERT

Red Hat CloudForms can use VMware alarms as a trigger for an alert. This type of alert can be created for a cluster, host, or virtual machine.

1. Navigate to **Control → Explorer**.

2. Click the **Alerts** accordion, then click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

3. In the **Info** area:

   - Type in a description for the alert.

   - From **Based On**, select **Cluster**, **Host**, or **VM**.

   - For **What to Evaluate**, select **VMware Alarm**.

   - In **Notification Frequency**, select how often you want to be notified if the log item is detected.

4. In the **VMware Alarm Parameters** area select the provider and alarm.

5. After setting the parameters, you then select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

6. Click **Add**.

6.9. CREATING AN EXPRESSION ALERT

Expression alerts enables you to create a notification based on any possible criteria for clusters, datastores, hosts, and virtual machines. The following procedure creates an alert for when a host's datastore has less than 5% free space.

1. Navigate to **Control → Explorer**.

2. Click on the **Alerts** accordion, then click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

3. In the **Info** area:

   - Type in a description for the alert.

   - From **Based On**, select **Host/Node**.

- For **What to Evaluate**, select **Expression (Custom)**.

- In **Notification Frequency**, select how often you want to be notified if the expression is evaluated to true.

4. Use the expression editor to create your expression. This is the same expression editor used to create **Conditions**. For details on how to use the expression editor, see the *Policies and Profiles Guide*.

Expression (Choose an element of the expression to edit)

5. Click ✔ (**Commit expression element changes**) to accept the expression.

6. After setting the parameters, you then select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

7. Click **Add**.

6.10. CREATE AN OPERATIONAL ALERT

1. Navigate to **Control → Explorer**.

2. Click on the **Alerts** accordion, then click ⚙ (**Configuration**), ⊕ (**Add a new Alert**).

3. In the **Info** area:

   - Type in a description for the alert.

   - Check **Active** when you feel that the alert is ready to be enabled.

   - From **Based On**, select **Server**.

   - Select the appropriate driving event.

   - In **Notification Frequency**, select how often you want to be notified if the event log threshold is reached.

4. After setting the parameters, select what you want the alert to do. You can send an email, create an SNMP Trap, or send a management event to start an automation process.

5. Click **Add**.

6.11. OPERATIONAL ALERT TYPES

Table 6.2. Operational Alerts

Driving Event	Explanation (Thresholds, Description)	Proposed Action and Next Steps
EVM Server Start	Alert is raised when an server starts.	Send e-mail. This is a notification.
EVM Server Stop	Alert is raised when a server stops.	Send e-mail. Review logs to see why the server stopped.
EVM Server Not Responding	Alert is raised when one server detects that another server has not responded in 2 minutes.	This is a sign of a problem that should be investigated. Check logs.
EVM Server Exceeded Memory Limit	Alert is raised when an server has exceeded its system memory limit and begins killing workers. Default is 80%. Threshold configured in Advanced Settings. server: :worker_monitor: :kill_algorithm: :name: :used_swap_percent_gt_value :value: 80	This may be caused by the following issues: The server is running with too few resources. The server is enabled with too many roles or number of workers. The server picked up all the roles because another server has failed. A runaway process has taken up most of the memory.
EVM Server is Master	When one server takes over as a master server.	Typically, this should only occur when first starting a set of servers, perhaps following expected outages. If a server picks up as master in other situations, the old master had an issue that needs to be researched (such as server not responding in time).

Driving Event	Explanation (Thresholds, Description)	Proposed Action and Next Steps
EVM Server High System Disk Usage	The server's system disk is 80% full. This check is run as part of a system schedule. Threshold configured in Advanced Settings. server: events: :disk_usage_gt_percent: 80	Temp files used by the operating system may be filling the disk. Yum updates, normal /tmp files, or temp files in /var/lib/data/miqtemp/ may cause the problem.
EVM Server High App Disk Usage	The server's app disk is 80% full. This check is run as part of a system schedule. Threshold configured in Advanced Settings. server: events: :disk_usage_gt_percent: 80	Server temp files may remain.
EVM Server High Log Disk Usage	The server's log disk is 80% full. This check is run as part of a system schedule. Threshold configured in Advanced Settings. server: events: :disk_usage_gt_percent: 80	Logs are getting too big or are not being log rotated properly every day. Review the most recent logs.
EVM Server High DB Disk Usage	The server's db disk is 80% full. This check is run as part of a system schedule. Applies if using PostgreSQL as the VDMB. Threshold configured in Advanced Settings. server: events: :disk_usage_gt_percent: 80	Database or database logging is getting too large. May require **full** vacuuming of PostgreSQL database.

Driving Event	Explanation (Thresholds, Description)	Proposed Action and Next Steps
EVM Worker Started	Alert is raised when a worker is about to start.	This is a notification. Failover may trigger this alert.
EVM Worker Stopped	Alert is raised when a worker is requested to stop.	Review logs for reason in the event the worker was not purposefully stopped.
EVM Worker Killed	Alert is raised when a non-responsive worker does not restart on its own and is killed.	Review logs for reason the worker was killed. May be the result of EVM Worker Not Responding.
EVM Worker Not Responding	Alert is raised when a worker has not responded for 2 minutes (:heartbeat_timeout) or has not started within 10 minutes (:starting_timeout).	An influx of events from the Virtual Center or Red Hat Virtualization causes an inability of EVM/CFME to handle the capacity at which they are being queued. Utilize the Event Handler Configuration to filter events that are causing problematic queue table growth.
EVM Worker Exceeded Memory Limit	Alert is raised when a worker exceeds the memory threshold. The default is 150 MB, but some workers have their own value in the :memory_threshold section for that specific worker.	Review logs for reason the worker is exceeding the memory limit. This may be the result of an overload to the worker process that requires further investigation.
EVM Worker Exceeded Uptime Limit	Alert is raised when a worker has been running longer than the :restart_interval. (Most workers are set to never restart using the 0.hours setting.) The EMS Refresh SmartProxy workers are set to restart every 2 hours.	Review logs for the reason the worker is exceeding the limit. This may be the result of an overload to the worker process that needs further investigation.
EVM Worker Exit File	Alert is raised when the scheduler worker exits due to a pending large ntp time change.	This is a notification.

6.12. EDITING AN ALERT

After creating an alert, you can edit the threshold, expression, or the notification type.

1. Navigate to **Control → Explorer**.

2. Click on the **Alerts** accordion, then click on the alert that you need to edit.

3. Click (**Configuration**), (**Edit this Alert**).

4. Make the required changes.

5. Click **Save**.

6.13. COPYING AN ALERT

You can copy an existing alert to create a new alert that is similar to the existing one, then change the values associated with it.

1. Navigate to **Control → Explorer**.

2. Click on the **Alert** accordion, then click on the alert that you want to copy.

3. Click (**Configuration**), (**Copy this Alert**). Click **OK** to confirm.

4. Make the required changes.

5. Click **Add**.
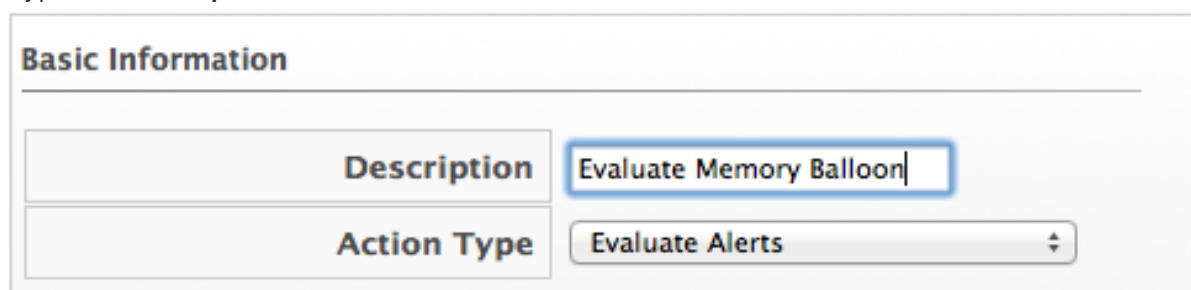
6.14. DELETING AN ALERT

When an alert is no longer needed, you can remove it from your VMDB.

1. Navigate to **Control → Explorer**.

2. Click on the **Alerts** accordion, then click on the alert that you want to delete.

3. Click (**Configuration**), (**Delete this Alert**).

4. Click **OK** to confirm.

6.15. EVALUATING AN ALERT

1. Navigate to **Control → Explorer**.

2. Click the **Actions** accordion, then click (**Configuration**), (**Add a new Action**).

3. Type in a **Description** for the action.

4. Select **Evaluate Alerts** from **Action Type**.

5. Select the alerts to be evaluated and click ▶ (**Move selected Alerts into this Action**). Use the **Ctrl** key to select multiple alerts.

6. Click **Add**.

CHAPTER 7. ALERT PROFILES

7.1. CREATING ALERT PROFILES

Alert profiles enable you to create groups of standard alerts. An alert profile can have as many alerts assigned as you need, and can be assigned to clusters, datastores, hosts, and virtual machines.

> **NOTE**
>
> Each profile type under the **Alert Profile** accordion menu will have its own **Alert Selections** and subsequent **Assignments** choices.

1. Navigate to **Control → Explorer**.

2. Click on the **Alert Profiles** accordion, then click on the type of profile that you want to create.

3. Click ⚙ (**Configuration**), ⊕ (**Add a new Profile**).

4. In the **Basic Information** box, type in a unique **Description** for the alert profile.

5. Select the desired alerts from the **Alert Selection** area. Use the **Ctrl** key to select multiple alerts.

6. Click ▶ to add the Alerts.

7. Type in any additional description in the **Notes** area.

8. Click **Add**.

7.2. EDITING AN ALERT PROFILE

You can edit an alert profile as your enterprise's needs change.

1. Navigate to **Control → Explorer**.

2. Click on the **Alert Profiles** accordion, then click the alert profile you want to edit.

3. Click ⚙ (**Configuration**), ✏ (**Edit this Alert Profile**).

4. Make the required changes.

5. Click **Save**.

7.3. DELETING AN ALERT PROFILE

Remove alert profiles that you no longer need. This does not remove the alerts associated with the alert profile.

1. Navigate to **Control → Explorer**.

2. Click on the **Alert Profiles** accordion, then click the alert profile you want to remove.

3. Click ⚙ (**Configuration**), 🗑 (**Delete this Alert Profile**).

4. Click **OK** to confirm.

7.4. ASSIGNING AN ALERT PROFILE

After an alert profile is created and verified, you can assign it directly to a resource.

1. Navigate to **Control → Explorer**.

2. Click on the **Alert Profiles** accordion, then click on the alert profile that you want to assign.

3. Click ⚙ (**Configuration**), ✏ (**Edit Assignments for this Alert Profile**).

4. The options presented change based on if the alert is for a cluster, datastore, Red Hat CloudForms server, host, or virtual machine and instance. You can assign to the enterprise, to specific hosts, cluster, resource pools, and providers, or based on assign tags. For a Red Hat CloudForms server alert profile, you can only assign to Red Hat CloudForms servers in the current Region.

5. Click **Save**.

CHAPTER 8. IMPORTING AND EXPORTING

8.1. IMPORTING AND EXPORTING POLICIES, POLICY PROFILES, AND ALERTS

If you have multiple VMDBs, you can export policies, policy profiles, or alerts from one to another. You can export and import for use with other Red Hat CloudForms infrastructures.

> **NOTE**
>
> The appliances you are exporting or importing data between must use the same CloudForms version.

1. Copy the file to import to a location that is accessible to your Red Hat CloudForms console.

2. Navigate to **Control → Import/Export**.

3. Click **Choose file** to navigate to the location of the file.

4. Select the file, and then click **Open** from the file selection box.

5. Click **Upload**.

6. Verify that these are the policies or policy profiles that you want to import.

7. Click **Commit**.

8.2. EXPORTING A POLICY, POLICY PROFILE, OR AN ALERT

1. Navigate to **Control → Import/Export**.

2. From the **Export dropdown**, select policy profiles, policies, or alerts, depending on what you want to export.
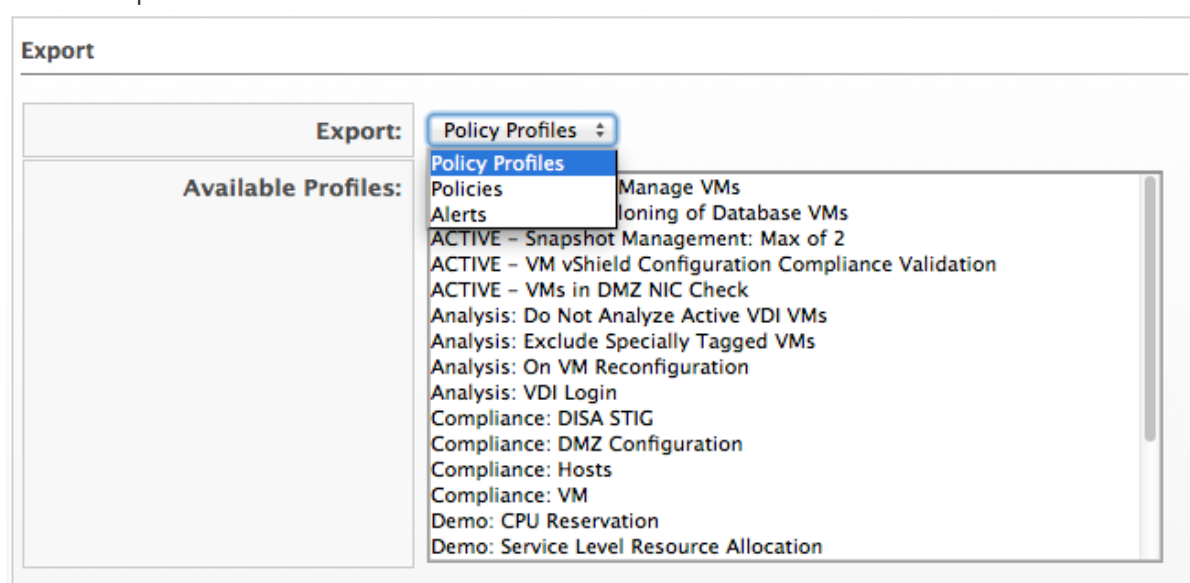
3. From the **Available Profiles** or **Available Policies** or **Available Alerts** list, select the items to export. Use the **Ctrl** key to select multiple items to export into one file.

4. Click **Export**.

5. Follow the prompts in your browser to save the file.

CHAPTER 9. OPENSHIFT PROMETHEUS ALERTS

> **NOTE**
>
> This feature is currently available as a technology preview only. For more information on the support scope for features marked as technology preview, see Technology Preview Features Support Scope.

Prometheus is used as an external alerting component. CloudForms integrates Prometheus alerts for use with OpenShift Container Platform. CloudForms collects events from Prometheus, generates alerts based on these events, and then attaches alerts to inventory objects.

It is possible to view ongoing alerts in CloudForms by navigating to **Monitor → Alerts** and manage their life cycle, including the ability to:

- View active alerts per provider in the **Overview** screen.

- View alert data and related objects.

- Comment, assign, and acknowledge alerts in the **All Alerts** screen.
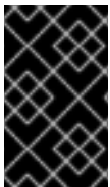
> **IMPORTANT**
>
> Network utilization metrics for OpenShift Container Platform are only collected for **eth0** network interfaces. Therefore, network utilization trend will not be displayed for network interfaces other than **eth0** on the overview screen of the provider.

9.1. CONFIGURING PROMETHEUS ALERTS IN CLOUDFORMS

Configuring and enabling Prometheus alerts in CloudForms comprises the following steps:

1. Deploying and configuring Prometheus on OpenShift.

2. Assigning Prometheus alert profiles to the enterprise in CloudForms.

3. Adding an OpenShift Containers Provider with a Prometheus alert endpoint in CloudForms.

9.1.1. Deploying and Configuring Prometheus on OpenShift

1. Deploy Prometheus on OpenShift by following the steps covered in Prometheus on OpenShift Container Platform.

2. Add Prometheus to an OpenShift cluster and configure alert definitions. See the example code block below for configuring alerts (currently must be done on the Prometheus side):

```
$ oc edit configmap -n openshift-metrics prometheus
# Supported labels:
# severity: ERROR|WARNING|INFO. defaults to ERROR.
# Supported annotations:
# description: The error message displayed on the screen
# miqTarget: ContainerNode|ExtManagementSystem, defaults to ContainerNode.
# miqIgnore: "true|false", should ManageIQ pick up this alert, defaults to true.
  prometheus.rules: |
    groups:
```

```
- name: example-rules
  interval: 30s # defaults to global interval
  rules:
  - alert: "NodeDown" << Alert starts here
    expr: up{job="kubernetes-nodes"} == 0
    annotations:
      miqTarget: "ContainerNode"
      severity: "ERROR"
      url: "https://www.example.com/node_down_fixing_instructions"
      description: "Node {{$labels.instance}} is down"
  - alert: "TooManyRequests"
    expr: rate(authenticated_user_requests[2m]) > 12
    annotations:
      miqTarget: "ExtManagementSystem"
      severity: "ERROR"
      url: "https://www.example.com/too_many_requests_fixing_instructions"
      description: "Too many authenticated requests"
```

3. Reload the Prometheus configuration. You can do this by deleting the Prometheus pod, then redeploying with the new configuration.

9.1.2. Assigning Prometheus Alert Profiles to the Enterprise

Complete the following procedure to assign Prometheus alert profiles to the enterprise using the CloudForms user interface.

> **NOTE**
>
> Both Node and Provider alert profiles are created automatically during the installation, so it is not required to create these profiles.

1. Navigate to **Control → Explorer**, then click **Alert Profiles** in the accordion menu.

2. Click to expand **Node Alert Profiles**, then click **Prometheus Node Profile**.

3. Click (**Configuration**), then click (**Edit assignments for this Alert Profile**). Assign the profile to **The Enterprise** from the **Assign To** list.

4. Click **Save**.

5. Click to expand **Provider Alert Profiles**, then click **Prometheus Provider Profile**.

6. Click **Configuration**, then click **Edit assignments for this Alert Profile**. Assign the profile to **The Enterprise** from the **Assign To** list.

7. Click **Save**.

> **NOTE**
>
> Alerts triggered before assigning the profiles to the enterprise will not appear in CloudForms at all.

9.1.3. Adding an OpenShift Containers Provider with a Prometheus Alert Endpoint

Complete the following procedure to add an OpenShift Containers Provider with a Prometheus Alert Endpoint using the CloudForms user interface.

1. Navigate to **Compute → Containers → Providers**.

2. Click (**Configuration**), then click (**Add a New Containers Provider**).

3. Enter a **Name** for the provider.

4. From the **Type** list, select **OpenShift Container Platform**.

5. Enter the appropriate **Zone** for the provider. If you do not specify a zone, it is set to **default**.

6. From the **Alerts** list, select **Prometheus** to enable external alerts. Selecting **Prometheus** adds an **Alerts** tab to the lower pane to configure the Prometheus service. Alerts are disabled by default.

7. From the **Metrics** list, select **Hawkular** or **Prometheus** to collect capacity and utilization data, or leave as **Disabled**. Selecting **Prometheus** or **Hawkular** adds a **Metrics** tab to the lower pane for further configuration. Metrics are disabled by default.

8. In the **Default** tab, configure the following for the OpenShift provider:

   a. Select a **Security Protocol** method to specify how to authenticate the provider:

   - **SSL**: Authenticate the provider securely using a trusted Certificate Authority. Select this option if the provider has a valid SSL certificate and it is signed by a trusted Certificate Authority. No further configuration is required for this option.

   - **SSL trusting custom CA**: Authenticate the provider with a self-signed certificate. For this option, copy your provider's CA certificate to the **Trusted CA Certificates** box in PEM format.

      NOTE

     You can obtain your OpenShift Container Platform provider's CA certificate for all endpoints (default, metrics, alerts) from **/etc/origin/master/ca.crt**. Paste the output (a block of text starting with **-----BEGIN CERTIFICATE-----**) into the **Trusted CA Certificates** field.

   - **SSL without validation**: Authenticate the provider insecurely (not recommended).

   b. Enter the **Hostname** (or IPv4 or IPv6 address) of the provider.

      IMPORTANT

     The **Hostname** must use a unique fully qualified domain name.

   c. Enter the **API Port** of the provider. The default port is **8443**.

   d. Enter a token for your provider in the **Token** box.

> **NOTE**
>
> To obtain a token for your provider, run the **oc get secret** command on your provider; see Obtaining an OpenShift Container Platform Management Token.
>
> For example:
>
> # oc get secret --namespace management-infra management-admin-token-8ixxs --template='{{index .data "ca.crt"}}' | base64 --decode

    e. Click **Validate** to confirm that Red Hat CloudForms can connect to the OpenShift Container Platform provider.

9. For the **Prometheus** alerts service, add the Prometheus alerts endpoint in the **Alerts** tab:

    a. Select a **Security Protocol** method to specify how to authenticate the service:

- **SSL**: Authenticate the provider securely using a trusted Certificate Authority. Select this option if the provider has a valid SSL certificate and it is signed by a trusted Certificate Authority. No further configuration is required for this option.

- **SSL trusting custom CA**: Authenticate the provider with a self-signed certificate. For this option, copy your provider's CA certificate to the **Trusted CA Certificates** box in PEM format.

- **SSL without validation**: Authenticate the provider insecurely using SSL. (Not recommended)

    b. Enter the **Hostname** (or IPv4 or IPv6 address) or alert **Route**.

    c. Enter the **API Port** if your Prometheus provider uses a non-standard port for access. The default port is **443**.

    d. Click **Validate** to confirm that CloudForms can connect to the alerts service.

10. If you selected a metrics service, configure the service details in the **Metrics** tab:

    a. Select a **Security Protocol** method to specify how to authenticate the service:

- **SSL**: Authenticate the provider securely using a trusted Certificate Authority. Select this option if the provider has a valid SSL certificate and it is signed by a trusted Certificate Authority. No further configuration is required for this option.

- **SSL trusting custom CA**: Authenticate the provider with a self-signed certificate. For this option, copy your provider's CA certificate to the **Trusted CA Certificates** box in PEM format.

> **NOTE**
>
> In OpenShift, the default deployment of the router generates certificates during installation, which can be used with the **SSL trusting custom CA** option. Connecting a Hawkular endpoint with this option requires the CA certificate that the cluster uses for service certificates, which is stored in **/etc/origin/master/service-signer.crt** on the first master in a cluster.

- **SSL without validation**: Authenticate the provider insecurely using SSL. (Not recommended)

   b. Enter the **Hostname** (or IPv4 or IPv6 address) of the provider, or use the **Detect** button to find the hostname.

   c. Enter the **API Port** if your Hawkular or Prometheus provider uses a non-standard port for access. The default port is **443**.

   d. Click **Validate** to confirm that Red Hat CloudForms can connect to the metrics endpoint.

11. Click the **Advanced** tab to add image inspector settings for scanning container images on your provider using OpenSCAP.

> **NOTE**
>
> - These settings control downloading the image inspector container image from the registry and obtaining the Common Vulnerabilities and Exposures (CVE) information (for effective scanning) via a proxy.
>
> - CVE URL that CloudForms requires to be open for OpenSCAP scanning: https://www.redhat.com/security/data/metrics/ds/. This information is based on the source code of OpenSCAP.

   a. Enter the proxy information for the provider in either **HTTP**, **HTTPS**, or **NO Proxy** depending on your environment.

   b. Enter the **Image-Inspector Repository** information. For example, **openshift3/image-inspector**.

   c. Enter the **Image-Inspector Registry** information. For example, **registry.access.redhat.com**.

   d. Enter the **Image-Inspector Tag** value. A tag is a mark used to differentiate images in a repository, typically by the application version stored in the image.

   e. Enter **https://www.redhat.com/security/data/metrics/ds/** in **CVE location**.

12. Click **Add**.

> **NOTE**
>
> You can also set global default image-inspector settings for all OpenShift providers in the advanced settings menu by editing the values under **ems_kubernetes**, instead of setting this for each provider.
>
> For example:
>
> ```
> :image_inspector_registry: registry.access.redhat.com
> :image_inspector_repository: openshift3/image-inspector
> ```

Once you have completed the procedure, you will have OpenShift Prometheus alerts enabled in CloudForms, and can manage their life cyle from the **Monitor → Alerts** screen in the user interface.

APPENDIX A. REGULAR EXPRESSIONS

In Red Hat CloudForms, regular expressions can be used to search the contents of a file for a specific string for use in a condition. Below are listed the items most commonly used with Red Hat CloudForms to search strings. These are a small subset of all the items available to use in regular expressions. If you are unfamiliar with regular expressions, there are many resources available on the Internet, including www.regular-expressions.info. Note that if you want to search a file, you must collect it as part of a host analysis profile.

Table A.1. Regular Expressions

Anchors	
^	start of string
$	end of string
Character Classes	
\s	white space including spaces, tabs, and line breaks
\S	not white space
\d	digit, same as [0-9]
\D	not digit
\w	word
\W	not word
Quantifiers	
*	0 or more of preceding characters
+	1 or more of preceding characters
?	0 or 1 of preceding character
Escape Character	
\	put before a metacharacter to search for that actual character
Metacharacters	
^ [. $ { * (\ +) | ? < >	must be used with the Escape Character if you are searching specifically for it

Special characters	
\n	new line
\t	tab
Groups and Ranges	
.	any character except new line (\n)
(a|b)	a or b
()	group
[abc]	a or b or c
[^abc]	not a or b or c
[a-q]	letter between a and q
[A-Q]	upper case letter between A and Q
[0-7]	digit between 0 and 7
Pattern modifiers	
i	case insensitive
Other helpers	
.*	swallows text between 2 words
\s+	guarantees minimum of 1 whitespace between 2 words
\s*	guarantees 0 or more whitespace between 2 words
^\s*	beginning of line with zero or more whitespace
\s+.*	swallows all text and white space between 2 words
\d+	guarantees minimum of 1 number between 2 words
<\w>	identical to <[a-zA-Z0-(_]>

Table A.2. Examples

Description	Regular Expression
([a-zA-Z0-9]+)	Letters, numbers, hyphens
Find the line beginning with **sshd**. Then, using a colon **:** as delimiter, check that the value four ":" over is equal to 99999.	^sshd:[^:]*:[^:]*:[^:]*:99999:
Verify that **PASS_MAX_DAYS** exists starting in position **1** and a value after it is ⇐ 90.	^\s*PASS_MAX_DAYS\s+([0-9]|[1-8][0-9]|90)
Verify that **ROOTPW** (in any case) exists on an uncommented line.	/[⌐#]*ROOTPW/i
Verify that line in file starts with size and the value after is ⇐ 4096k.	^\s*size\s+(409[0-6]|40[0-8][0-9]|[123][09]{3}|\d{1,3})k
Find line with string **restrict 127.0.0.1** that starts in position `1 to ensure it is not commented out.	^\s*restrict\s+127\.0\.0\.1
Find an uncommented line that contains /**home**. There will be additional text before the desired string.	[⌐#]*\/home

APPENDIX B. REPORTABLE FIELDS IN RED HAT CLOUDFORMS

Red Hat CloudForms provides a set of out of the box reports covering a variety of areas. These fields are collected by Red Hat CloudForms from the different parts of the infrastructure. Most of the fields are named directly as the vendor would name them. However, there are some fields that Red Hat CloudForms calculates given this data. Note that this list may vary by Provider and may change as new Providers are added to Red Hat CloudForms. Also, note that the value returned by each reportable field below is in megabytes (MB).

Table B.1. Clusters

Reportable Field
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max

Reportable Field
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg

Table B.2. Datastore Files

Reportable Field
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
VM Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
VM Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
VM Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
VM Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VM and Instance : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VM and Instance : Memory – Avg Used for Collected Intervals 30 Day Avg
VM and Instance : Memory – Avg Used for Collected Intervals 30 Day High Avg
VM and Instance : Memory – Avg Used for Collected Intervals 30 Day Low Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg

Reportable Field
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Logical Disk.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
File Share.File System.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
File Share.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg

Reportable Field
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg

Table B.3. Datastores

Reportable Field
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.File System.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.File System.Logical Disk.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.File System.Logical Disk.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg

Reportable Field
File Share.File System.Logical Disk.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
File Share.File System.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
File Share.File System.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
File Share.Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
File Share.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
File Share.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
File Share.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
File Share.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Hosts : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg

Reportable Field
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg

Table B.4. EVM Groups

Reportable Field
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max

Reportable Field
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg

Table B.5. Hosts

Reportable Field
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Miq Templates : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max

Reportable Field
VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg

Table B.6. Images

Reportable Field
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Cluster : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Datastore : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Datastores : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Host : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Host : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Host : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg

Reportable Field
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Provider : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Provisioned From Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Provisioned From Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Provisioned From Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Provisioned From Template : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Provisioned VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Avg
Provisioned VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day High Avg
Provisioned VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Low Avg
Provisioned VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Service.VMs : Memory – Aggregate Used for Child VMs for Collected Intervals 30 Day Max
Service.VMs : Memory – Avg Used for Collected Intervals 30 Day Avg
Service.VMs : Memory – Avg Used for Collected Intervals 30 Day High Avg
Service.VMs : Memory – Avg Used for Collected Intervals 30 Day Low Avg