

vRealize Hyperic Monitoring and Management Guide

vRealize Hyperic 5.8.4



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2014 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 About vRealize Hyperic Monitoring and Management Guide 5**
- 2 Getting to Know the vRealize Hyperic Dashboard 6**
 - Portlet Quick Facts and Links 7
 - Customize the Dashboard 8
 - Auto-discovery Portlet 9
 - Understanding the Availability Summary Portlet 11
 - Working with the Control Actions Portlet 12
 - Favorite Resources Portlet 13
 - Group Alert Summary Portlet 14
 - Metric Viewer Portlet 14
 - Problem Resources Portlet 15
 - Recent Alerts Portlet 16
 - Saved Charts Portlet 17
 - Summary Counts Portlet 18
- 3 Monitor Resources in the Resource Tab 19**
 - Browsing to Resources 20
 - Indicators Minitab 21
 - Metric Data Minitab 22
 - Resources Minitab 25
 - Metric Display Range 29
 - Charting Metric Data 30
 - Reading a Full-page Chart 32
 - Map Control 35
 - Metric Extrapolation View for Groups 35
 - View Metrics for Top Processes 36
- 4 Global Management Views 38**
 - View Alerts in the Alert Center 38
 - View Alert Definitions in the Alert Center 40
 - Event Center 41
 - Viewing Currently Down Resources 42
 - Working with the Operations Center 43
 - Managing Reports 44
- 5 Views for Specific Resource Types 49**
 - GemFire View 49

- [HQ vSphere View](#) 50
- [tc Server Application Management View](#) 78
- [Agent Control View](#) 83
- [Spring Insight View](#) 86

6 [Running Live System and Metric Queries](#) 92

- [Creating JMX MBean Queries for JVMs](#) 92
- [Running Live Exec Queries](#) 94

7 [Using Resource Control Actions](#) 96

- [Running Resource Control Actions](#) 96
- [Viewing Control Action Status and History](#) 98

8 [View Alerts and Update Alert Status](#) 100

- [Listing Recent Alerts](#) 100
- [List Alerts for a Resource](#) 100
- [Viewing Alerts Across All Resources](#) 101
- [Monitoring Unfixed Alerts](#) 101
- [View Alert Details](#) 101
- [Acknowledging or Fixing an Alert](#) 101

About vRealize Hyperic Monitoring and Management Guide

1

The *vRealize Hyperic Monitoring and Management Guide* describes how to use the VMware[®] vRealize Hyperic[®] user interface, perform resource control actions, and process fired alerts.

Intended Audience

vRealize Hyperic Monitoring and Management is intended for operations personnel who monitor resources with vRealize Hyperic.

Getting to Know the vRealize Hyperic Dashboard

2

The vRealize Hyperic dashboard is the first page displayed in the Hyperic user interface when you log on. This topic describes its various components.

The vRealize Hyperic dashboard is the first page displayed in the vRealize Hyperic user interface when you log on. The dashboard contains multiple portlets, each of which presents a particular type of information, such as resource health, recent alerts, recently discovered resource or changes to resources, recently performed resource control actions, and so on.

You can personalize the dashboard to suit your responsibilities and preferences; you can add or remove portlets, rearrange them, and customize the contents of many of the portlets.

In vRealize Hyperic, you can have multiple dashboards - one for each role to which you are assigned. For more information, see .

Note also that in vRealize Hyperic, the roles you are assigned to govern what resource data appears in dashboard portlets. Portlets will contain only the resource data to which your Hyperic roles allows access.

- [Portlet Quick Facts and Links](#)

This topic lists the different portlets available in the Hyperic Dashboard, and their descriptions.

- [Customize the Dashboard](#)

This topic describes how to customize the vRealize Hyperic dashboard.

- [Auto-discovery Portlet](#)

The Auto-Discovery portlet on the vRealize Hyperic Dashboard lists recently added or modified platforms and servers, and allows an authorized user to view discovery details, and to import the new or changed resource data to the database.

- [Understanding the Availability Summary Portlet](#)

The Availability Summary portlet presents the availability of selected resources by resource type. The portlet is empty until you configure the resources to include.

- [Working with the Control Actions Portlet](#)

The Control Actions portlet displays information about recent control actions performed on resources over a configurable period of recent history.

- [Favorite Resources Portlet](#)

The Favorite Resources portlet lists the name, resource type, current availability, and total fired alerts for selected resources.

- [Group Alert Summary Portlet](#)

The Group Alerts Summary portlet displays traffic light indicators for resource alerts and group alerts for selected groups.

- [Metric Viewer Portlet](#)

The Metric Viewer portlet displays a selected metric for selected resources of the same resource type.

- [Problem Resources Portlet](#)

The Problem Resources portlet lists resources that have recently fired an alert or out-of-bounds metric.

- [Recent Alerts Portlet](#)

The Recent Alerts portlet presents a list of recently fired alerts.

- [Saved Charts Portlet](#)

The Saved Charts portlet allows you to view charts of interest from the dashboard.

- [Summary Counts Portlet](#)

The Summary Counts portlet on the HQ Dashboard shows resource counts by inventory type, and if configured to do so, counts for selected resource types.

Portlet Quick Facts and Links

This topic lists the different portlets available in the Hyperic Dashboard, and their descriptions.

Table 2-1. Portlet Quick Facts and Links

Portlet	Description	More Information and Configuration Options
Autodiscovery	Lists new and changed resources, and allows you to add them to inventory.	
Availability Summary	Displays the availability of selected of resources, aggregated by resource type. Refreshes every minute. This portlet does not contain any resources until you configure it.	
Control Actions	Lists the control actions that vRealize Hyperic has most recently performed on any managed resource, upcoming scheduled control actions, and the most frequently performed quick control actions.	
Favorite Resources	Lists resources that have been added to the portlet. This portlet does not contain any resources until you configure it.	
Saved Charts	Displays charts that have been saved to the Dashboard. In vRealize Hyperic, charts are displayed as a slideshow.	

Table 2-1. Portlet Quick Facts and Links (Continued)

Portlet	Description	More Information and Configuration Options
Recent Alerts	Lists the most recently triggered alerts for a set of user-selected resources. Refreshes every minute.	
Recently Added	Lists platforms recently added to inventory.	
Search Resources Portlet	Allows you to search for resources. The search supports case-insensitive, partial-term queries and is restricted to a single inventory type.	
Summary Counts	Displays a count of all the managed resource of each inventory type. In vRealize Hyperic, the resources that appear are limited to those that you have permissions to view.	
Group Alert Summary	Displays traffic light indicators for resource alerts and group alerts for selected groups. Click on the traffic light for a group alert to view a list of alerts that have fired. Click on a group name to view the group page. This portlet does not present any groups until you add them to it.	
Metric Viewer	Shows selected metrics for selected resources. Refreshes every minute. This portlet does not contain any resources until they have been added to it.	
Problem Resources	Lists all resources with problem metrics, their availability, the number of alerts triggered on each, the number of times the metric has been out of bounds, and the last time the OOB metric was collected. Click the resource name to view the Monitor page for it.	

Customize the Dashboard

This topic describes how to customize the vRealize Hyperic dashboard.

You can change a portlet's layout and behavior.

Procedure

- 1 Grab a portlet by the bar at the top of it and drag it to the desired location.
- 2 Click the gear-shaped icon in the bar at the top display a configuration dialog.

Add Portlets

This topic describes how to add a portlet.

Procedure

- ◆ At the bottom of each column of a portlet, click **Add Content** to this column to list portlets you can add.

Remove Portlets

You can remove portlets that are no longer useful or of interest.

To remove a portlet from the dashboard:

Procedure

- ◆ Click the Delete icon in the top right corner of the portlet.

You can add a removed portlet at a later time. For more information, see [Add Portlets](#).

Auto-discovery Portlet

The Auto-Discovery portlet on the vRealize Hyperic Dashboard lists recently added or modified platforms and servers, and allows an authorized user to view discovery details, and to import the new or changed resource data to the database.

By default, the Auto-Discovery portlet lists the (up to) five most recently new or changed platforms.

A platform appears in the Auto-Discovery portlet if it, or a server running on it, is new or changed. The rules are these:

- The platform is new — A platform appears as new in the portlet, if neither its IP address or FQDN match that of an existing platform in inventory. In this case, any new servers discovered on the platform appear below the platform in the portlet.
- The platform has a new server — A new server has been discovered on the platform since the last scan. The new server is listed below the platform.
- Platform or server properties have changed — One or more inventory properties for the platform, or for one or more of the servers running it, have changed since the last scan. Servers with changed properties are listed below the platform.

The Auto-Discovery portlet presents the following information for each platform it contains:

- **Hostname** — The hostname of the platform is a link to a page — the Auto-Discovery Results page - that contains detailed information about the scan results for the platform and servers running there. See the *vRealize Hyperic Administration Guide* for more information about using the Auto-Discovery page.
- **Platform type** — The resource type for the platform.

- **Status** — Indicates the type of change that was detected for a resource, either "new" or "modified".
- **Changes** — If the Status for a resource is "modified", the Changes column contains a summary of what changed. For example:
 - "server set changed" — Applies to platforms; this value indicates that changes to one or more servers on the platform were detected. The changed server(s) are listed below the platform.
 - "name change" — Indicates that the name of the resource has changed; a resource name can change when a resource is upgraded from one version to another, if version number forms a portion of the resource name, as is often the case.
 - "install path changed" — Indicates that the installation path for a server has changed; the installation path for a resource can change when a resource is upgraded from one version to another, if version number forms a portion of the path, as is often the case.
 - "IP set changed" — Indicates that the IP address has changed. When the agent detects an IP address not associated with an existing platform in inventory, it checks for a platform with a matching FQDN - if found, Hyperic recognizes the platform as existing.
 - "FQDN changed"
- For each newly discovered or changed server on the platform:
 - Installation path
 - Status - Indicates new or modified
 - Changes - If status is modified, the **Changes** column contains a summary of what changed.

Import or Skip Resources in Auto-Discovery Portlet

When the vRealize Hyperic agent discovers all of the resource properties required to monitor a resource, it starts monitoring that resource as soon as you add it to inventory. This is the case for most resource types.

However, some level of configuration is required to start managing some resources types. You can process the contents of the Auto-Discovery Portlet in these ways:

- To import all resources---leave all resources selected, and click **Add to Inventory**.
- To skip all resources---leave all resources selected, and click **Skip Checked Resources**.

Procedure

- 1 To import all resources, leave all resources selected and click **Add to Inventory**.
- 2 To skip all resources, leave all resources selected, and click **Skip Checked Resources**.
- 3 To import selected resources, either:
 - Clear the checkboxes for the resources that you do not want to add to inventory and click **Add to Inventory**.
 - Clear the checkboxes for the resources you do want to add to inventory and click **Skip Checked Resources**.

When you do not import a resource displayed in the Auto-Discovery portlet:

- You skip its servers, too.
- During the next platform scan, skipped resources will reappear in the portlet after the next scan that detects them. If you have resources that you do not want the agent to discover, see the relevant section in the *vRealize Hyperic Administration Guide*.

Configure the Number of Auto-Discoveries Displayed

To configure the number of auto-discoveries that are displayed in the portlet, follow these instructions.

Procedure

- 1 Click the gear icon in the upper left corner of the portlet.
- 2 On the Display Settings page, select "10" or "all", and click **OK**.

Understanding the Availability Summary Portlet

The Availability Summary portlet presents the availability of selected resources by resource type. The portlet is empty until you configure the resources to include.

Configure the Availability Summary Portlet

Follow this process to configure the Availability Summary portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Availability Summary portlet.
The HQ Dashboard Settings: Availability Summary page appears.
- 2 In the **Description** field, enter the title to appear at the top of the portlet.
- 3 In the **Display Range** pulldown, select the maximum number of resource types to list in the portlet.
- 4 If the Selected Resources section contains resources that you do not want to be included in the availability summary, select the respective checkbox and click **Remove From List**.
- 5 Click **Add to List** to add resources to the portlet.
- 6 In the Resources column of the Dashboard Settings: Availability Summary Add/Remove Resources page:
 - a Use the View filters to filter by inventory type, resource type, or both, as desired.
 - b Enter a substring in the **Filter by Name** field to filter by resource name, as desired.

- c In the list of resources, checkmark desired resources, and click the right arrow between the columns to include them in the availability summary.
- d If the Add Resources column contains resources you do not wish to include in the availability summary, select the checkboxes for the respective resources and click the left arrow between the panels.

7 Click **OK**.

Working with the Control Actions Portlet

The Control Actions portlet displays information about recent control actions performed on resources over a configurable period of recent history.

The Control Actions portlet displays information about recent control actions performed on resources over a configurable period of recent history.

- **Recent Control Actions** - This section lists the individual resource control action performed during the configured period, starting with the most recent. The following information is displayed for each resource:
 - **Resource Name** - Name of the resource.
 - **Control Action** - The control action that was performed.
 - **Date/Time** - When the action was performed.
 - **Message** - Results of the control action.
- **Quick Control Frequency** - This section lists the resources upon which the most on-demand control actions have been performed during the configured interval. The following information is displayed for each resource:
 - **Resource Name** - Name of the resource.
 - **# of Control Actions** - How many control actions were performed during the interval.
 - **Most Frequent Control Action** - The control action that was most frequently performed.

Configure the Control Actions Portlet

This topic provides instruction for configuring the Control Actions portlet.

To configure the amount of data in the Control Actions portlet:

Procedure

- 1 Click the gear icon in the upper right corner of the Control Actions portlet.
The HQ Dashboard Settings: Control Actions page appears.

- 2 Configure what appears in the Recent Control Actions section using the first row of filters.
 - a Check the box next to the **Control Action Range** if you want to list recent control actions in the portlet. (To stop the display of this data, uncheck it.)
 - b In the first drop-down list, select the maximum number of completed control actions to list.
 - c In the second drop-down list, select the time range from which to gather that data.
- 3 Configure what appears in the Quick Control Frequency section using the second row of filters:
 - a Check the box if you want to list the resources with the most on-demand control actions. (To stop the display of on-demand control-action data, uncheck it.)
 - b Select the maximum number of resources to list.
- 4 Click **OK**.

Favorite Resources Portlet

The Favorite Resources portlet lists the name, resource type, current availability, and total fired alerts for selected resources.

Select a Favorite Resource

This topic describes how to select a favorite resource.

You can add a resource to Favorite Resources when you are viewing it in the Resources tab.

Procedure

- ◆ In the Resources tab, click **Add to Dashboard Favorites** on the **Tools** menu.

Configure Favorite Resources on the Dashboard

This topic describes how to configure your favorite resources on the dashboard.

Procedure

- 1 Click the gear icon in the upper right corner of the Favorite Resources portlet.
The HQ Dashboard Settings: Favorite Resources page appears.
- 2 If the Selected Resources section contains resources that you do not want to be included in the portlet, select the respective checkbox and click **Remove From List**.
- 3 Click Add to List to add resources to the portlet.
The Dashboard Settings: Favorite Resources: Add Resources page appears.

- 4 In the Resources column:
 - a Use the View filters to filter by inventory type, resource type, or both, as desired.
 - b Enter a substring in the Filter by Name field to filter by resource name, as desired.
 - c In the list of resources, checkmark desired resources, and click the right arrow between the columns to include them in the portlet.
- 5 If the Add Resources column contains resources you do not wish to include in the portlet, checkmark those resources and click the left arrow between the panels.
- 6 Click **OK**.

Group Alert Summary Portlet

The Group Alerts Summary portlet displays traffic light indicators for resource alerts and group alerts for selected groups.

Click on the traffic light for a group alert to view a list of alerts that have fired. Click on a group name to view the group page. This portlet does not present any groups until you add them to it.

Configure the Group Alerts Portlet

This topic provides instructions for configuring the Group Alerts portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Group Alerts Summary portlet.

The portlet refreshes.

 - The Available groups column lists compatible groups that have not been added to the portlet.
 - The Enabled groups column lists compatible groups that have been added to the portlet.
- 2 Change the content of the portlet by clicking a group, and using the appropriate arrow to move the group from one column to another.
- 3 Click **Save**.

Metric Viewer Portlet

The Metric Viewer portlet displays a selected metric for selected resources of the same resource type.

Configure the Metric Viewer Portlet

This topic provides instructions for configuring the Metric Viewer portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Metric Viewer portlet.

The HQ Dashboard Settings: Metric Viewer page appears.

- 2 Specify the display settings for the portlet:
 - Description — This description shows up at the top of the portlet.
 - Display Range — The total number of resources to display in the portlet.
 - Resource Type — This selection filters both the set of resources that users can add to the portlet and the list of metrics available for selection.
 - Metric — The single metric that will be displayed for the user-selected set of resource in the portlet.
 - Sort Order — Determines whether the resources will be displayed in the portlet by descending or ascending metric value/
- 3 If the Selected Resources section contains resources that you do not want to be included in the portlet, select the respective checkbox and click **Remove From List**.
- 4 Click Add to List to add resources to the portlet.
The Dashboard Settings: Metric Viewer: Add/Remove Resources page appears.
- 5 In the Resources column:
 - a Enter a substring in the **Filter by Name** field to filter by resource name, as desired.
 - b In the list of resources, checkmark desired resources, and click the right arrow between the columns to include them in the portlet.
- 6 If the Add Resources column contains resources you do not wish to include in the portlet, select the checkbox for the respective resources and click the left arrow between the panels.
- 7 Click **OK**.

Problem Resources Portlet

The Problem Resources portlet lists resources that have recently fired an alert or out-of-bounds metric.

The Problem Resources portlet on the HQ Dashboard lists resources that, over a configured period of recent history, have either had a fired alert, or an out-of-bounds metric. An out-of-bounds metric is a metric that had a value outside the "acceptable" range of values set by the baselining process.

Note Tracking out-of-bounds metrics is a behavior that is configured globally, in the "Automatic Baseline Configuration Properties" section of the Administration > HQ Server Settings page. If tracking of OOB metrics is not enabled, OOB metrics will not cause a resource to appear in the Problem Resources portlet.

The following information is shown for each problem resource:

- Current availability of the resource.
- Number of alerts that have fired during the historical period.

- Number of out-of-bounds metrics reported during the historical period.

Resources are ordered by inventory type (platform, server, and service), and then by the date and time of problem occurrence, from most recent to least recent.

Configure the Problem Resources Portlet

This topic provides instructions for configuring the Problem Resources portlet.

You can use the Problem Resources page to limit the resources that are included in the Problem Resources portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Problem Resources portlet.
The HQ Dashboard Settings: Problem Resources page appears.
- 2 To limit the resources that appear in the portlet:
 - Show Maximum of — Select maximum number of problem resources to list.
 - For the Last — Choose the period of history for which to show problem resources.
- 3 Click **OK**.

Recent Alerts Portlet

The Recent Alerts portlet presents a list of recently fired alerts.

The following information is shown for each alert:

- Date/Time - When the alert fired. Click to view the Alert Detail page.
- Alert Name - Name of the alert definition.
- Resource Name - The resource where the alert fired.
- Fixed - Whether the alert has been marked "fixed".
- Ack - whether the alert has been acknowledged.

Configure the Recent Alerts Portlet

Follow this process to configure the Recent Alerts Portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Recent Alerts portlet.
The HQ Dashboard Settings: Recent Alerts page appears.
- 2 In the Display Settings section:
 - Description — (optional) Enter a description, if desired. The description will appear at the top of the portlet.

- Alert Range — Select:
 - Maximum length of the alert list: 4, 10, 20 or 30.
 - The minimum priority of the alerts to include.
 - Timeframe of interest; choices range from 30 minutes to a month.
 - Resources of interest:
 - all resources - With this setting, alerts (that meet the priority and timeframe criteria) on any resource can appear in the list.
 - selected resources - With this setting, the alert list contains only resources listed in the Selected Resources section that meet the priority and timeframe criteria.
- 3 If desired, define a set of resources of interest in the Selected Resources section.
- The content of the portlet will be limited to alerts that fired on resources in the Selected Resources only when the **all resources/selected resources** switch is set to **selected resources**. The contents of the Selected Resources section is ignored when the switch is set to **selected resources**.
- 4 In the Resources column:
- a Use the View filters to filter by inventory type, resource type, or both, as desired.
 - b Enter a substring in the Filter by Name field to filter by resource name, as desired.
 - c In the list of resources, checkmark desired resources, and click the right arrow between the columns to include them in the portlet.
- 5 If the Add Resources column a contains resources that are not of interest, select the respective checkbox of those resources and click the left arrow between the columns.
- 6 Click **OK**.

Saved Charts Portlet

The Saved Charts portlet allows you to view charts of interest from the dashboard.

The charts are shown as a slideshow.

Add Chart to Saved Charts Portlet

This topic provides instructions for adding a chart to your dashboard.

Procedure

- ◆ When viewing a chart, click **Save to Dashboard**.

Remove a Chart from Saved Charts Portlet

This topic provides instructions for removing a chart from your dashboard.

Procedure

- ◆ Click the **Remove Chart** button on the chart.

Summary Counts Portlet

The Summary Counts portlet on the HQ Dashboard shows resource counts by inventory type, and if configured to do so, counts for selected resource types.

By default, the Summary Counts portlet shows the total resources of each inventory type — Application, Platform, Server, Service, Compatible Group, and Mixed Group — as shown in the screenshot below.

Configure the Summary Counts Portlet

Follow this process to configure the Summary Counts Portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the portlet.
The HQ Dashboard Settings: Summary Counts page appears.
- 2 To exclude the count for a particular inventory type, (for instance, Applications) clear the **Show Total of All ...** checkbox for the inventory type.
- 3 To show the count of resources of particular resource types, (for instance, WebLogic Admin 8.1 and MySQL 5.x), check the box next to the respective resource types.
- 4 Click **OK**.

Monitor Resources in the Resource Tab

3

These topics describe how to use the Resources tab component of the VMware vRealize Hyperic user interface.

- [Browsing to Resources](#)
This topic describes the Browse page.
- [Indicators Minitab](#)
This section discusses the various types charts and graphs available in the Indicators minitab.
- [Metric Data Minitab](#)
The Metric Data minitab, available when an individual resource, or a group of resources of the same type are selected, is a tabular summary of the measurements collected during the currently selected metric display range.
- [Resources Minitab](#)
The Resources minitab includes the following information.
- [Metric Display Range](#)
By default, the last 8 hours of history are presented on the Indicators page and other views in the Monitor tab. There are two ways to change the display range.
- [Charting Metric Data](#)
This topic provides information about the metric data charts.
- [Reading a Full-page Chart](#)
This topic provides information about understanding full-page charts.
- [Map Control](#)
This topic provides information about the Map control.
- [Metric Extrapolation View for Groups](#)
This page allows you to predict future values of a metric for the members of group based on historical metrics.
- [View Metrics for Top Processes](#)
You can view the top processes that are consuming resources at a point in time to identify the cause of a problem on a specific platform.

Browsing to Resources

This topic describes the Browse page.

The first time you display the **Resources > Browse** page during a session, it lists the platforms in inventory, and the type, description, and current availability for each. To list resources of another inventory type, click the link for that type — **Servers**, **Services**, **Compatible Groups/Clusters**, **Mixed Groups**, or **Applications** — above the resource list.

Listing Resources by Resource Type

This topic discusses the different resource views.

After you select an inventory type to list, a resource type selector list appears. The pull-down label depends on the currently selected inventory type, for example, **All Platform Types**, **All Server Types**, and so on. Select a platform type from the list to list only resources of that type.

If the currently selected inventory level is **Servers** or **Services**, a pull-down list of resource types of the selected inventory type is present. Select a resource type from the list to list only resources of that type.

If the currently selected inventory level is **Compatible Groups/Clusters**, a pull-down list labelled "Group Type" will appear. The list will contain resource types. If there is at least one compatible group of a resource type, that type will appear in the pull-down. Select a resource type from the list to list only groups that contain that type.

If the currently selected inventory level is **Mixed Groups**, a pull-down list labelled "Group Type" will appear. The entries in the list depend on the mixed groups configured in your environment. Select a group type from the list to list only groups of that composition. Depending on your deployment, the list may include:

- Mixed Group - Platforms, Servers and Services
- Mixed Group - Groups
- Mixed Group - Applications

Filtering by Groups

Follow this process to filter platforms, servers, and services by groups.

When the currently selected inventory level is **Platforms**, **Servers** or **Services**, a pull-down list of groups appears. It lists both compatible and mixed groups to which you have access.

Procedure

- ◆ To list the platforms, servers, or services that are members of a particular group, select the group from the **All Groups** drop-down menu.

Filter Resources by Additional Characteristics

This topic provides instructions for filtering by name, owner, or availability.

You can apply additional filters the list of resources currently displayed.

Procedure

- 1 To filter by resource name, enter an alphanumeric string that matches all or a portion of resource name.
- 2 Click the **Unavailable** box to filter the current list of resources to include only resources that are currently unavailable.
- 3 Click the **Owned by > UserName** box to filter to include only resources
- 4 Click the green control to the right of the filter options to apply the filters.

Indicators Minitab

This section discusses the various types charts and graphs available in the Indicators minitab.

The **Indicators** minitab contains a chart for each indicator metric for the currently selected resource.

Availability Bar and Timeslices

This topic describes the availability bar.

The bars across the top and bottom of the the tab contain grey dots, each of which corresponds to a timeslice. The length of the timeslices depends upon the currently selected **Metric Display Range** — if the display range is 8 hours, each timeslice is 8 minutes. If the display range is 4 hours, each timeslice is 4 minutes. The color of the dot for a timeslice indicates the availability of the resource during that timeslice. If the timeslice is longer than the collection interval for the metric, the availability shown for the timeslice is based on the multiple data points collected during the timeslice.

- Green — indicates 100% availability during timeslice — each time that availability was reported during the interval, the resource was available.
- Orange — indicates availability greater than 0% and less than 100% — during the interval, the availability of the resource was reported to be unavailable data point during the interval was the resource was available.
- Red — indicates 0% availability during timeslice — each time that availability was reported during the interval, the resource was not available.

The average availability over the display range is shown to the right of the availability bar.

Click an availability indicator to display the start time of the timeslice, and a vertical bar to help you view the state of each metric during the timeslice.

Metric Charts in the Indicators Minitab

This section describes the metric charts and how they are displayed.

The **Indicators** page displays a chart of each of the resource's indicator metrics. The values labeled **LOW**, **AVG**, and **PEAK** are the lowest, average, and highest values collected during the metric display range.

The indicators are displayed as column charts: the area of each column indicates the value range (the high and low values) of the metric. The average value of the metric is indicated by the cross in the column. The charts are stacked vertically so that their X-axis (time) values line up. This works in conjunction with the vertical highlight from the availability and timeline bars to analyze the metric data across multiple metrics for a specific time. This is useful when trying to diagnose a problem at a specific time by correlating relevant metric values. Arrangements of charts for selected metrics can be saved as a "view" so that users can easily always compare the same metric data and therefore understand the interaction between different resources.

- To change the display order of the charts, click the up and down arrow controls at the upper-right-hand corner of each chart.
- To remove a chart from display, click the red X icon at the upper-right-hand corner of the chart.
- To save the set of displayed charts and their order, either
 - Update the current view: Select "Update <current-view name>" in View and click the icon.
 - Create a new view: Select "Create New View" in View, type a view name, and click icon!
- To delete the current view: Select "Delete View" in View and click icon.png!
- To display another view: Select "Go to View" in View and select one of the views.

Event Icons and Popups

This topic describes the different event icons and popups.

A purple circular icon over a timeslice in the bar at the bottom of the Indicators pane indicates an event was logged occurred during that timeslice. Click an event icon to display details in a popup.

Metric Data Minitab

The Metric Data minitab, available when an individual resource, or a group of resources of the same type are selected, is a tabular summary of the measurements collected during the currently selected metric display range.

By default, the Metric Data minitab displays only metrics for which measurements exist during the metric display range. To list all metrics supported for the resource, click the control next to Show All Metrics at the top of the Metric Data minitab.

When all supported metrics for a resource are listed, a Hide Metrics Without Data control replaces the Show All Metrics control - click the the control next to it to toggle the display.

The columns in the Metric Data minitab vary, depending on whether an individual resource or a group of resources is selected, as described in the subsections below.

Metric Data Minitab for a Single Resource

The Metric Data minitab displays the following data for each metric that is enabled for one or more members of a compatible, over the current metric display range:

- Alerts — Number of times a collected metric value triggered an alert.
- OOB — Number of times the metric was out-of-bounds.
- LOW — Lowest value collected.
- AVG — Average of values collected.
- PEAK — Highest value collected.
- LAST — Last collected value.
- Collection Interval — Frequency of metric collection. "NONE" indicates that data is not being collected.

Metric Data Minitab for an Autogroup or Compatible Group

The Metric Data minitab displays the following data for each metric that is enabled for one or more members of a group of resources of the same type — an autogroup or a compatible group.

- Number Coll — Number of data points collected across all group members for which the metric is enabled.
- Alerts — Number of times a collected metric value triggered an alert.
- OOB — Number of times the metric was out-of-bounds across all group members.
- LOW — Lowest value collected across all group members.
- AVG — Average of values collected across all group members.
- PEAK — Highest value collected across all group members.
- SUM — Sum of all values collected, not for a metric whose unit is percentage.
- Collection Interval — Frequency of metric collection.
 - "NONE" — indicates that data is not being collected.
 - "VARIES" — Indicates that the collection interval varies among members of the group.
- Member health data — (For compatible groups only, not autogroups) The lower portion of the Metric Data minitab for a compatible group lists the resources in the group, the current availability of each, and in the rightmost column. an icon that links to the Alerts page for the resource.

Compare Metrics for Resources in a Compatible Group

Check the member resources in "Current Health of group's <group name> resources collecting metrics" and click **Compare Metrics of Selected** at the bottom of the page.

Proceed to the Compare Metrics screen.

View Metric Details and Metadata

To display metric data for a single resource in a separate window, click the log at the right end of metric data for the resource.

Tailoring Metric Collection Settings for the Resource

You can tailor the metric collection settings for the resource on the Metric Data minitab.

The default metric collection settings for a resource are specified on the Administration > Monitoring Defaults page for the associated resource type. You can alter the metric collection interval for a specific resource on its Metric Data minitab — note however that subsequent updates to the monitoring defaults for the resource type will overwrite any modifications to collection intervals made for a specific resource. For information about setting metric collection options for all resources of the same resource type, see the Tailor Metric Collection for a Resource Type page of the *vRealize HypericAdministration Guide*.

Changes to make to metric collection settings for a group will apply to all members of that group.

Configuring Metric Baselines

You can use the Set Baselines control at the bottom of the Metric Data minitab to reset the baseline, acceptable high, and acceptable low values for a metric.

Enable Collection of a Metric

This topic provides instructions for enabling the collection of a metric.

Procedure

- 1 1. Click the arrow to the right of **Show All Metrics** at the top of the minitab.
All metrics, including currently disabled metrics, will be listed.
- 2 Select the checkbox for each metric you want to enable.
- 3 Specify the frequency of the metric collection:
 - a a. Enter an integer value in the **Collection Interval for Selected** field.
 - b Select **Minutes** or **Hours** from the pull-down list.
 - c Click the arrow to the right of the **Minutes/Hours** pull-down list.

If the currently selected resource is a compatible group, collection of the metric is enabled for all resources in the group.

Disable Collection of a Metric

This topic provides instructions for disabling a metric.

Procedure

- 1 Select the checkbox alongside each metric you wish to disable.
- 2 Click **Disable Collection** at the bottom of the minitab.

If the currently selected resource is a compatible group, collection of the metric is disabled for all resources in the group.

Set Collection Interval for a Metric

This topic provides instructions for setting the collection interval for a metric.

Procedure

- 1 If the Metric Data minitab currently display all metrics for the resource, including disabled metrics, click **Hide Metrics Without Data**.
- 2 Select the checkbox alongside each metric whose collection interval you wish to modify.
You can checkmark more than one metric if you want to set the same collection interval for each of them.
- 3 To specify the frequency of metric collection:
 - a Enter an integer value in the **Collection Interval for Selected** field.
 - b Select **Minutes** or **Hours** from the pull-down list.
 - c Click the arrow to the right of the **Minutes/Hours** pull-down list.

Resources Minitab

The Resources minitab includes the following information.

- [Resource Lists in the Resources Minitab](#)
This topic lists the resources in the Resources minitab.
- [Availability Icons in the Resources Minitab](#)
The color of the icon in the **Availability** column indicates the current availability status of the resource, or the resource group.
- [Availability Details in the Resources Minitab](#)
Hover over the callout icon for more information about the resource. The information displayed in the popup varies by resource type.
- [View Metric Data for Related Resources in the Resources Minitab](#)
This topic provides instructions for viewing metric data for related resources.

- [Use Controls and Popup Options in the Resources Minitab](#)

This topic provides instructions for viewing additional information about each metric.

- [Filter Metrics by Category and Value Type](#)

This topic provides filtering options for filtering metric data.

Resource Lists in the Resources Minitab

This topic lists the resources in the Resources minitab.

Table 3-1. Resource Lists in the Resources Minitab

For a ...	the Resource minitab lists these related resources....	
Platform	Platform Services Health Deployed Servers Health	<p>List of platform services running on the platform</p> <p>List of servers running on the platform. If more than one service or server of the same type runs on a host, they are presented as a single item in the appropriate list - an autogroup - whose name is the resource type.</p>
Server	Services Host Platform	<p>List of services running in the server. If more than one service runs on a server, they are presented as a single item - an autogroup - whose name is the resource type.</p> <p>Platform where the server runs.</p>
Service	Host Server	Server where the service runs.
Compatible Group	Host Platforms Host Servers	<p>Group Members - List of resources in the group, and either:</p> <p>If the compatible group contains servers, this section lists each platform upon which one or more group member runs.</p> <p>If the compatible group contains services, this section lists each server upon which a group member runs.</p>
Applications	Services Host Servers	<p>List of the services that the server comprises. If the application contains more than one service of the same type, they are presented as a single item in the Services list - an autogroup - whose name is the resource type.</p> <p>Lists each server upon which a service in the applications runs.</p>

Availability Icons in the Resources Minitab

The color of the icon in the **Availability** column indicates the current availability status of the resource, or the resource group.

Table 3-2. Availability Icons in the Resources Minitab

Availability Icon Color	State	Definition
Green	Up	For an individual resource, indicates that its availability status is "up". For a group, indicates that none of the group members has availability status of "down".
Yellow	Warn	This state is valid for a group only. It indicates that (1) one or more group members, but not all members, have the availability state of "down", and (2) the remaining members have status "green".
Blue	Maintenance/Paused	This icon appears only for a VM or a group of VMs. For an individual VM indicates it is paused. For a group, indicates that all VMs in the group are paused.
Red	Down	For an individual resource, indicates that its availability status is "down". For a group, indicates that all members of the group have either availability status "red" or "unknown".
Grey	Unknown	For an individual resource, indicates that its availability is unknown. For a group, indicates that the availability of at least one of the members is unknown.
Blue	Suspended/Powered off	This icon appears only for a VM or a group of VMs. For an single VM indicates it is suspended or powered off. For a group of VMs, indicates that all VMs in the group are suspended or powered down.

Availability Details in the Resources Minitab

Hover over the callout icon for more information about the resource. The information displayed in the popup varies by resource type.

Table 3-3. Availability Details in the Resources Minitab

For a ...	The Popup Contains
Platform, Platform Service, Server, or Service	<ul style="list-style-type: none"> ■ Inventory level and resource type of the resource. ■ Last known throughput metric value for the resource.
Autogroup	<ul style="list-style-type: none"> ■ Inventory level and resource type of the resource. ■ Breakdown of resources availability by availability status.
Compatible Group	<ul style="list-style-type: none"> ■ Resource type for the resources in the group. ■ Last known throughput metric value for the group. ■ • Go to Resources link — Click to expand the group and see a list of its members.

View Metric Data for Related Resources in the Resources Minitab

This topic provides instructions for viewing metric data for related resources.

Procedure

- 1 Select the checkbox next to each resource for which you wish to display metrics.
- 2 Filter metrics as desired:
 - Problem Metrics — Only metrics whose values are outside the expected range established via the baselining process will be listed. For each metric, the portlet displays two counts: the number of the selected resources that have "OOB" values for the metric, and the number of "Alerts" that the metric value has triggered
 - All Metrics — The indicator metrics for selected resources will be listed - regardless of whether they are out of bounds.
- 3 Click **View Metrics**.

Use Controls and Popup Options in the Resources Minitab

This topic provides instructions for viewing additional information about each metric.

When you use the View Metrics option to list metrics, two controls are available for each metric, to the right of the **Alerts** column.

Procedure

- ◆ Hover over the callout icon next to a metric to display a popup menu of options.
 - Click **Chart Metric in Indicators** to add a chart for the metric to the chart pane to the right.
 - Click **View Full Chart** to view a full page chart for the metric.

- Click **Metric Data** to display the Metric Details and Metadata for the metric in a new window.

Filter Metrics by Category and Value Type

This topic provides filtering options for filtering metric data.

When the Metric Data minitab is active, rather than the Indicators minitab, the Resources minitab provides filtering options.

Procedure

- 1 Uncheck a metric category (for example, Availability, Throughput, Utilization, or Performance), to exclude metrics of that type from the list of metrics.
- 2 Uncheck a value type (for example, Dynamic, Trends Up, Trends Down, or Static) to exclude metrics with that value type from the list of metrics.
- 3 Click the green arrow control to filter the list.

Metric Display Range

By default, the last 8 hours of history are presented on the Indicators page and other views in the Monitor tab. There are two ways to change the display range.

Select a Number of Minutes, Hours, or Days

Use the pull-downs in the Metric Display Range section of the page to select a numeric value (available choices range from 4 to 120) and a unit of time (minutes, hours or days), and click the green right-arrow to refresh the page.

Define Advanced Settings for Metric Display Range

This topic provides instructions for defining advanced settings to display the metric range.

For more control over the display range:

Procedure

- ◆ Click **Advanced Settings** in the Metric Display Range section of the page. On the popup that appears, either:
 - Click **Last**, enter a number, and select a unit of time (minutes, hours or days) and click REDRAW, or
 - Click **Within a Date Range**, and use the pull-downs and fields to designate the start and end of the desired display interval, and click REDRAW.

After the page is redrawn, the Metric Display Range section shows the currently selected range. You can use the Show Last 8 Hours to return to the default display range, or Edit Range to display the Advanced Settings popup and choose a different display range.

Charting Metric Data

This topic provides information about the metric data charts.

A metric chart is a graphical representation of the metric data that vRealize Hyperic collects for each resource. Each resource type supports a different set of metric data, so charts vary by resource type also. Charts can help you see trends in resource behavior, which can in turn be helpful for capacity planning, cost analysis, and other operational decision-making.

vRealize Hyperic displays all full-page metric charts on the Metric Chart page, where you can adjust how the metric is displayed. (Indicator charts are more minimally displayed on the Current Health screen.) Consult the help for those screens for the details about viewing and interpreting charts.

How a Chart is Constructed

A metric chart by default displays only the measured metric data values over a user-selected time frame.

A metric chart by default displays only the measured metric data values over a user-selected time frame. You can also chart peak (highest observed), lowest and highest acceptable, and average values, and you can change the time frame that the chart covers. If you are charting a single metric for a single resource, you can also establish and chart baseline values. Depending on the display range over which you are charting the metric data and whether the metric data is cumulative (for example, throughput), the values may be averaged over time.

Chart Types

This topic provides information about the various chart types and how to view them.

What They Are and How to View Them

There are three kinds of charts, and each gives you a different perspective on resource health. For any chart, you have to select the metrics to be charted: this selection occurs on the Current Health screen for the resource (or for the parent group if looking at multiple resources), on the **Monitor** tab, on the **Metric Data** tab, which lists all the metrics collected for the resource being viewed. The chart of the selected metric(s) is displayed on the Metric Chart screen.

Table 3-4. Chart Types

Type	Description
Single Metric Single Resource	<p>This chart shows a single metric for a single resource. To chart a single metric:</p> <ol style="list-style-type: none"> 1 On the Metric Data tab, click the name of the metric or next to the metric whose chart you want to see. <p>Another way to view the chart page for a metric for a single resource is from the Indicators tab on the monitor visibility pages. If an indicator chart represents a metric for a single resource, clicking on its name will also display the full chart page for that metric and resource.</p>
Multiple Metrics Single Resource	<p>This chart shows up to a maximum of ten metrics simultaneously for a single resource, which is handy for getting an overall sense of the resource's health. To chart multiple metrics for a single resource:</p> <ol style="list-style-type: none"> 1 On the Metric Data tab, check the names of the metrics you want to chart and click Chart Selected Metrics. <p>The Indicators tab also displays multiple metrics for a single resource.</p>
Single Metric Multiple (Group) Resources	<p>This chart shows a single metric for a maximum of ten resources, which allows you to compare the behavior of different resources. The resources must be members of a single compatible group or of a single autogroup.</p> <ul style="list-style-type: none"> ■ • To chart a single metric for multiple resources of a compatible group: <ol style="list-style-type: none"> a On the Metric Data tab, click the name of the metric or next to the metric whose chart you want to see. <p>A full-page chart is displayed with data from the first 10 resources listed on the Metric Data tab.</p> b To change the set of resources displayed in the chart, on the Metric Chart screen, select or clear the checkboxes for the desired resources in Participating Resources and click Redraw Selected on Chart. <p>The chart will be refreshed to reflect metric data from the new set of resources.</p> ■ • To chart a metric for multiple members of an auto group: <ol style="list-style-type: none"> a Navigate to the parent resource that contains the auto group. b From the parent resource's Current Health screen, click the subtype. The Current Health screen for the auto group of that subtype is displayed. c On the Metric Data tab, click the name of the metric whose chart you want to see.

To view the chart for an indicator metric:

Instead of the **Metric Data** tab, on the **Indicators** tab, click the name of the indicator metric above the chart. A full-page chart is displayed.

Reading a Full-page Chart

This topic provides information about understanding full-page charts.

Each bar on a chart represents a timeslice that is 1/60 of the selected metric display range. A chart can plot a maximum of 60 points. The value of the Y-axis at the top of the bar is the observed metric value for that timeslice. What the value represents varies by the nature of the metric:

Static For a static metric, one whose value does not vary over time (for instance a timestamp), the plotted value corresponds to a single measurement.

Dynamic For a dynamic metric, one whose value may go up or down over time:

- The plotted point is the average of the values collected over the timeslice that ended at the time shown for the bar in the X-axis. The length of time over which the value was averaged depends on the current metric display range. For a display range of eight hours, each charted point represents the average of the preceding eight-minute period (8 hours / 60 time slices along the X-axis = 8 minutes). If the metric is collected every 60 seconds, and the chart's display range is 60 minutes, each of the 60 plotted points on the graph represents the single, observed value for the metric at a single point in time.
- The I-bar superimposed on the vertical bar shows the range of values collected during the timeslice.

• Trends Up and Trends Down For a cumulative metric, one whose values either always increases or decreases, (such as bytes served, uptime, minimum response time), the plotted point is the maximum (or minimum) value for the time slice. The value shown for a metric that for trends up or down metrics is not the average over the timeslice.

More About Timeslices

Charts are divided into a maximum of 60 slices, where each slice represents the same amount of time, over any display range the user chooses. For example:

- If the charted metric's collection interval is 5 minutes and the selected metric display range is one hour, during which the metric value was collected 12 times, the chart will contain 12 time slices, and each slice will represent a single observed metric value.
- If the metric display range is 5 hours, during which the metric value was collected 60 times, the chart will display 60 time slices, each will represent a single observed metric value.

- If the selected display range is 10 hours long, during which the metric value was collected 120 times, the chart will display 60 timeslices - the maximum - and each bar will summarize the 2 data points collected during the timeslice. For dynamic metrics, data points are averaged. For trends up or down metrics the highest/lowest value is plotted.

Peak, Average, and Low Values

By default, three colored horizontal lines appear on a metric chart that show the peak (pink), average (green) and low (blue) values collected for the metric for duration plotted.

Baseline and Acceptable High and Low Values

In vRealize Hyperic, this additional information may appear on a chart:

- Baseline — The baseline value for the metric is shown as a tan horizontal line.
- Low Range — A pink horizontal bar appears across the chart, between the low range value and the maximum value shown on the chart's Y-axis.
- High Range — A green horizontal bar appears across the chart, between the high range value and the minimum value shown on the chart's Y-axis.

Participating Resources

This topic describes the contents of the Participating Resources section of the full-page chart.

This section lists the resource and its metrics or the resources (if you are charting metrics for a group) that are plotted in the chart. Users can change the individual resource's metrics or the set of group member resources that are included in the chart.

Metric Baseline and Expected Range

This topic provides information about the Metric Baseline and Expected Range section of the full-page chart.

This section allows the user to view and change the values current calculated or set for the metric baseline and the acceptable range of metric values (outside of which observed metric values will be considered out-of-bounds (OOB). This section is available only when a single metric is charted.

Filtering Chart Data

This topic provides information about the types of filters that you can apply.

The chart legend allows users to select types of data to include in the chart:

- Actual: The observed metric value at any point in time
- Peak: The highest observed metric value (this will be a horizontal line)
- Average: The average of observed metric values across the entire graph (this will be a horizontal line)
- Low: The lowest observed metric value (this will be a horizontal line)

- Low Range: The user-specific lowest acceptable metric value (this will be a horizontal line)
- High Range: The user-specific highest acceptable metric value (this will be a horizontal line)
- Baseline: The HQ- or user-set baseline value (this will be a horizontal line)
- Control Actions: Indicates on the chart when control actions were performed on the resource (to help correlate the actions with changes in performance)

Change the Metric Display Range

This topic provides instructions for changing the time range that the chart spans.

At the bottom right of the "Metric Chart" section

Procedure

- 1 To move forward or back eight hours, click the back-arrow or the forward-arrow, respectively.

The new display range is indicated.

- 2 For a specific date/time range:

- a Click **Edit Range**.
- b For a range counting back from the current time, select a length of time.
- c For a date/time range: Select a date/time range.
- d Click **Redraw**.

The metric charts on the screen automatically refresh to reflect the new display range. This display range value applies to all resource-monitoring screens.

Save the Chart

This topic provides instructions for saving the chart to different formats.

Procedure

- 1 To save the chart to the dashboard, at the top of the page, click **Save Chart to My Dashboard** to add the chart to the Saved Charts portlet in the dashboard.

If you have edit permissions for multiple dashboards, the Tools menu will instead have a Save Chart to Dashboards option. Select the option to view a list of dashboards for which you have edit permissions. You can select one or more dashboards, and click **Add** to add the chart to those dashboards.

- 2 To save the chart to a CSV format, at the top of the page, click **Export to CSV**.

Map Control

This topic provides information about the Map control.

The Map control, available when a resource is selected in the user interface, presents a graphical view of the current resource and resources that are related to it. The map illustrates hierarchical inventory relationships, and a resource's membership in groups or applications.

A resource map shows the name and resource type of currently selected resource in orange font. Child or member resources are shown above the selected resources; parent or containing resource is shown below.

In the example below,

- The name of the currently selected resource is in orange font. In the example, the resource is a server type: "ActiveMQ Embedded 5.3".
- The resources listed above the current resource are its children --- assuming it has children. If the currently selected resource is a service type, it would not have children.
- The resource below the current resource is its parent --- assuming it has a parent. If the currently selected resource is a platform type, it would have no parent.

Each resource name is a link; click the link to navigate to the Indicators page for that resource.

The resource type for each resource is shown below the resource name.

You can map a:

- platform
- server
- service
- compatible group

Metric Extrapolation View for Groups

This page allows you to predict future values of a metric for the members of group based on historical metrics.

The Metric Data Extrapolation plugin is available on the Views tab when a group is selected. Although this function is available for any group, whether compatible or mixed, metric projections are most meaningful for compatible groups.

The extrapolation method is a simple linear regression using the Apache Commons Math package. The historical metric data for a group is used to calculate a slope, which is used to predict future values.

To extrapolate a metric for a members of a group:

- 1 Navigate to compatible group of interest.
- 2 Choose the Metric Data Extrapolation option from the Views tab.

3 You can select:

- Metric to extrapolate. By default, Availability is selected. If desired, choose another metric.
- Known Data Range. The amount of history upon which the extrapolation will be based. You can select values ranging from 1 day to 1 year. Typically, a longer data range yields a better extrapolation. In some cases, however, it might be better to select the data range based on the timing of environment-related events, such as a new release or infrastructure change.
- Projection Range. The total duration included in the trend analysis, including the Known Data Range and the Projection Range. The range you select determines, but is not equal to, how far into the future the metric will be projected.
- Threshold value for the metric. Charts will display a horizontal line at the metric value you enter.

4 Click Update to plot the extrapolated values for the selected metric for each member of the Group. Historical values are plotted in blue.

Understanding Extrapolation Charts

This topic provides information about extrapolation charts.

Each chart has the same scale, to better show differences between the resources, and ease analysis. This is different from other HQ charts which are scaled independently. The heading for each chart shows:

- The resource name and metric charted
- The units of the metric charted
- Extrapolation confidence; values can be 'Excellent', 'Good', 'Average', 'Questionable', or 'Poor'.

Extrapolated values are represented by the orange triangle that appears between "now" and the end of the projection range on the timeline.

View Metrics for Top Processes

You can view the top processes that are consuming resources at a point in time to identify the cause of a problem on a specific platform.

You view the metrics for the top processes by clicking a time point on the Platform Indicator chart for a specific platform. The Top Processes popup displays metrics for CPU usage, Free Memory and Load Average resources. The available metrics are the same on each tab, but are ordered differently according to the resource type. The listed processes might differ between tabs because the top processes might differ between the resources.

You enable the collection on the **Resources** tab.

Procedure

- 1 On the **Resources > Platforms > Inventory** tab, click **Edit** to open the Configure Properties dialog box.
- 2 Under Monitoring, verify that the **top_processes.enable** checkbox is selected.

- 3 (Optional) Specify the number of processes for which the metrics are collected in the **top_processes.number** text box.

The default is ten processes.

- 4 Specify the interval between the metric collections in the **top_processes.interval** text box.

The default interval is one minute.

- 5 Click **Ok** to save your settings.

The values that you specified appear on the **Resources** tab under Configuration Properties.

When you click a time point on a Platform Indicator chart, the Top Processes popup appears. If the platform was only recently discovered, so has not been running for very long, the processes might not appear. You can change the metric display range to 60 minutes to see the top processes over the last minutes.

Global Management Views

These topics describe pages in the vRealize Hyperic user interface that provide deployment-wide views of resource status.

- [View Alerts in the Alert Center](#)
This topic provides instructions for viewing alerts in the Alert Center.
- [View Alert Definitions in the Alert Center](#)
This topic provides instructions for viewing alert definitions.
- [Event Center](#)
This topic describes the Event Center.
- [Viewing Currently Down Resources](#)
The Currently Down Resources page, available on the **Resource** tab is lists resources that are not currently available, up to a maximum of 1000 resources.
- [Working with the Operations Center](#)
This topic describes the various parts of the Operations Center.
- [Managing Reports](#)
This topic provides information about working with reports.

View Alerts in the Alert Center

This topic provides instructions for viewing alerts in the Alert Center.

The Alert Center page is a deployment-wide view of alerts and alert definitions

Procedure

- 1 Click the **Analyze** tab

The Alert Center page appears, with the **Alerts** tab selected.

The Alerts page displays the following data for each fired alert that matches the currently selected filter criteria, ordered chronologically.

- Date — The date and time the alert was triggered.
- Alert Definition — The name of the alert definition that prompted the alert.

- Resource+ — The resource the alert was triggered on.
- Fixed — Whether or not the alert was fixed.
- Acked by — The name of the user who acknowledged the alert.
- Priority — The alert's priority.

2 Click a column heading to sort the table by the content of that column.

Filtering the Alerts List

This topic provides a list of options for filtering the list of alerts.

These are the options for filtering the list of alerts, and their default settings:

Table 4-1. Filter Alert List

Option	Description	Values	Default Setting
Show	Filters the list of alerts by state.	<ul style="list-style-type: none"> ▪ Not Fixed ▪ In Escalation ▪ All 	All
Alert type	Filters the list of alerts by the type of target resource.	<ul style="list-style-type: none"> ▪ Resource — With this setting, only alerts fired on individual resources will appear. This includes alerts applied at the resource type level. ▪ Group — With this setting, only alerts that fired on resource groups appear. 	All
Minimum Priority	Filters the list of alerts by the priority of the associated alert definition.	<ul style="list-style-type: none"> ▪ Low ▪ Medium ▪ High 	Low
In the last	Limits the list to alerts fired with a recent period of time.	<ul style="list-style-type: none"> ▪ All Time ▪ Day ▪ 2 days ▪ 3 days ▪ 4 days ▪ 5 days ▪ 6 days ▪ Week 	All Time
Groups	Limits the list of alerts to those fired on resources in a group.	The pulldown lists every resource group (whether compatible or mixed) in the deployment, and the options "All Groups".	All Groups

Fix or Acknowledge an Alert in the Alert Center

This topic provides instructions for acknowledging a fix or escalation.

To mark alerts fixed and acknowledge an alert in escalation

Procedure

- ◆ Select the checkbox for the desired alerts, and click **Fixed** or **Acknowledge** at the bottom of the page.

View Alert Definitions in the Alert Center

This topic provides instructions for viewing alert definitions.

The **Definitions** tab on the Alert Center page lists alert definitions in your deployment — resource alerts, resource type alerts, and group alerts — that match the current filter criteria.

Procedure

- 1 Click the **Analyze** tab.

The Alert Center page appears, with the **Alerts** tab selected.

- 2 Click the **Definitions** tab to display the Resource Alert Definitions pane.

The pane displays the following data for each alert definition that matches the currently selected filter criteria, ordered by creation date, with most recently created definitions first.

- Name — The alert definition's name. Click to view the alert definition in the Resource Hub.
- Created — Date and time the alert was defined
- Modified — Date and time the alert was modified (if at all)
- Active — Whether or not the alert definition is configured to generate alerts) or not. If the alert is defined to fire once, and not again until after it is fixed, a yellow flag icon will be displayed in this column if the alert is not fixed.
- Last Alert — Date and time the alert was last triggered. Click to view a list of alerts that fired for the alert definition.
- Resource — The resource, resource type, or group for which the alert was defined. Click to view the Monitor tab for the resource in the Resource Hub.
- Escalation — The name of the escalation scheme assigned to the alert definition. Click to view the escalation.
- Priority — The alert definition priority.

- 3 Click a column heading to sort the table by the content of that column.

Filter Alert Definition List

This topic provides information about filters available for the definition list.

These are the options for filtering the list of alerts, and their default settings:

Table 4-2. Filter Alert Definition List

Option	Description
Definition Type	Filters the list of alert definitions by whether the definition was for an individual resource, a group of resources, or a resource type target — so you can limit the list to definitions assigned to a resource, or to a resource type. Values are: <ul style="list-style-type: none"> ■ Resource — To list all alert definitions that assigned to individual resources or to a resource type. ■ Group — To list only alert definitions defined for a resource group. ■ Resource Type — To list only alert definitions that are assigned to resource groups. Definition Type
Exclude Type-Based Definitions	This option is available only when "Resource" is selected for "Definition Type" — checkmark to exclude resource type alert definitions from the list.
Show Only Disabled Definitions	Further limit the definitions listed to those that are currently disabled.

Event Center

This topic describes the Event Center.

The Event Center page, available on the **Analyze** tab in the Masthead, is a deployment wide view of events that have been logged for resources in your vRealize Hyperic deployment.

vRealize Hyperic logs alerts as events automatically. You can configure vRealize Hyperic to log events for log messages and resource configuration changes.

Events are ordered by date, with most recent events first. To reverse the sort order, click the control in the Date column header.

Filter the Event List

You can filter the events that are displayed in the Event List.

To limit the events displayed in the Events section

Procedure

- ◆ Choose one or more of the following filter options
 - Minimum Status — Select a status to limit the log tracking events in the list to those of that level or higher. Log tracking event levels, in decreasing order of severity, are:
 - Error
 - Warning
 - Info

- Debug
- Type — Select a type to limit events in the list to those of that type:
 - Log track
 - Event Track
 - Alerts
- Time Range — Select on the following time ranges to limit the events in the list to those occurred during the range:
 - Last 4 hours
 - Last 8 hours
 - Last day
 - Last week
 - Last month

Groups — Select one or more groups to limit the events in the list to those related to resources in the selected group(s).

Contents of the Event Center Page

This section details the contents of the Event Center page.

The following data is shown for each event that matches the filter criteria selected in the Filter pane. This information is shown for each event in the list:

- Date — Date and time the event happened
- Status — For log track events, the level is shown. For alerts, the status column contains "alert".
- Resource — The name of the resource that caused the event
- Subject —
 - For an alert, the alert name.
 - For a configuration tracking event, the type of change: "add", "delete", "modify", or "rename".
- Detail — Information about the event detail or triggering condition.

Viewing Currently Down Resources

The Currently Down Resources page, available on the **Resource** tab lists resources that are not currently available, up to a maximum of 1000 resources.

The right pane of the Currently Down Resources pane lists resources that match the filter criteria selected in the Resource Types panel on the left, ordered by Down Time, from longest to shortest. Click a column header to sort by that column's contents — except for the "Alerts" column.

By default, the page will list up to 50 resources. You can choose to list up to 100, or up to 1000, using the **Show Most Recent** control.

The following data is shown for each resource on the list:

- Resource — Name of the down resource. Click to view the resource's Indicators page.
- Type — The resource's platform, server, or service type.
- Down Since — The time at which the resource became unavailable.
- Down Time — The length of time the resource has been unavailable.
- Alerts — Click the icon in this column to see a list of alerts fired for the resource.

Filter Resources of the Currently Down Page

Follow this process to view all currently unavailable resources of a particular type.

To list all currently unavailable resources of a particular inventory level:

Procedure

- 1 click one of the top level links in Resource Types pane.
 - Platforms
 - Servers
 - Services
- 2 To further narrow the list to unavailable resources of a particular resource type, expand the link for the inventory level, and select the desired resource type from list that appears.

Working with the Operations Center

This topic describes the various parts of the Operations Center.

The Operations Center page, available on the **Analyze** tab in the Masthead lists resources that are down, have unfixed alerts, or both. You can use filters to scope the content of the page to the resources and problem type of interest. This concise view includes shows the current number of unavailable resources and unfixed alerts, and a single line problem summary for each resource.

Use the Table Controls to specify the number of resources listed per page and the frequency with which the page content is updated.

Contents of Operations Center Page

This topic provides information about the contents of the Operations Center page.

The Operation Center page contains summary and detailed information for the alert or availability events that match the current Display Filter settings.

Current Filter Totals

The tables in the Current Filter Totals section summarize the results that match the current filter settings:

- Resources table - Current number of unavailable platforms and the current number of unavailable resources of any type - platform, server, or service. If the current Status Type selection limits the page to alerts, the count column contains N/A.
- Alerts table - Total number of unfixed alerts and number of unfixed alerts in escalation, broken down by alert priority. If the current Status Type selection limits the page to alert events, the cells contain N/A.

Resource Details for All Hosts

Each row in the table in the lower part of the Operations Center page represents a resource that is currently down or a resource that has an unfixed alert. If a resource is currently down and has an unfixed alert, it will appear in two rows.

- Platform - The platform where the resource runs.
- Resource - Name of the resource that is unavailable or has an unfixed alert.
- Alert Name - In a row for a resource with an alert, the name of the alert definition that fired the alert.
- Priority - In a row for a resource with an alert, the alert priority.
- Status Type - Indicates whether the row corresponds to a "Resource Down" event or an unfixed "Alert".
- Last Escalation - In a row for a resource with an alert in escalation, timestamp for the last escalation step performed.
- Last Check - For an unavailable resource, shows when metrics for the resource were last collected. For an alert, shows when it last fired.
- Duration - Indicates how long ago the alert fired or the resource went down.
- State - For an alert, a megaphone icon indicates the alert is in escalation. A man-at-work icon indicates the alert in escalation has been acknowledged.
- Status Information - For an alert, shows how many times it has fired and the condition that triggered it. If a metric condition triggered the alert, the triggering and current values of the metric are shown.

Managing Reports

This topic provides information about working with reports.

The Report Center has built-in reports that provide visibility into availability, alerts, inventory, resource utilization, and resources that are not collecting metrics. You can also create your own reports. Report Center reports can accept parameter values, and be generated in PDF, HTML, Excel, and CSV formats.

The Report Center is available under the **Analyze** tab.

Only a user with the Super User role may run reports in the Report Center.

Built-in Reports

This topic lists the reports available in the Report Center.

The table below describes the ready-to-run reports in the Report Center. vRealize Hyperic provides the following report templates:

Table 4-3. Built-in Reports

Report Template	What Does it Tell Me	What The Report Displays	Required User Input (Parameters)
Alert Detail	Which resources in my environment are firing the most alerts?	The same as the Alert Summary Report, except, instead of a summary of alerts, every alert is shown. This is useful when analyzing a full alert log.	<ul style="list-style-type: none"> ■ Start Date ■ End Date ■ Group
Alert Summary	Which resources in my environment are firing the most alerts?	A summary of all alerts that were triggered during the user-specified time range. You can restrict the list of alerts by a user-specified group. For each resource, the number of low-, medium-, and high-priority alerts are shown, along with the total number of alerts.	<ul style="list-style-type: none"> ■ Start Date ■ End Date ■ Group
Availability Downtime	How many minutes were resources unavailable over the given time range?	Availability percentage and minutes of downtime, over a user-specified time range, for resources in a user-specified group. This report cannot be run against the entire environment.	<ul style="list-style-type: none"> ■ Start Date ■ End Date ■ Group
Disk Usage	Which disks are running low on free space? Which disks are filling up the quickest?	Disks in the environment that are running low on free space, over a user-specified time range. This report can be run against a single group or across the entire environment.	<ul style="list-style-type: none"> ■ Start Date ■ End Date ■ Group
Inventory	What resources were created during the given time range? What new services have been discovered? (Hyperic auto-imports services and so they are not immediately obvious.)	All resources created in the Hyperic environment in the user-specified time range, sorted by creation date. This report operates on the entire environment.	<ul style="list-style-type: none"> ■ Start Date ■ End Date

Table 4-3. Built-in Reports (Continued)

Report Template	What Does it Tell Me	What The Report Displays	Required User Input (Parameters)
Event Log Summary	What's happening on each of my resources? A summary of events across all resources, instead of for a single resource.	Resources and their logged events, taken from the system's event logs, sorted by the number of events.	<ul style="list-style-type: none"> ■ Start Date ■ End Date
License Count	The platforms for which I have Hyperic licenses.	Platform name, type, license count, and date added.	N/A
Metric Interval	How many metrics are being collected on my resources, how often, and when? Basically, the report shows the metric load.	The collection interval and the number of metrics that are being collected for each resource.	N/A
Network Interface	What kind of traffic are my resources experiencing?	Bits in/out per second for each resource	<ul style="list-style-type: none"> ■ Start Date ■ End Date ■ Group
Resources Not Collecting Metrics	For which of my resources are metrics no longer being reported, even though they're supposed to be?	Resources that have stopped publishing metrics, plus the last time a metric was recorded for the resource	N/A
Resources Without Metrics	Which of my resources don't have metric collection enabled?	Resources that don't have metric collection enabled in Hyperic	N/A
Unconfigured Resources	Which of my resources are not configured for management? For instance, Oracle needs a password, so when it is discovered and imported, metrics will not be collected until the password is provided.		N/A

For reports involving specific resource types (for example, the Disk Usage and Network Usage reports include only File Server Mount resources), only compatible groups of that resource type are allowed as an input parameter. For all non-resource-type-specific reports (for example, the Alert and Availability reports), all groups defined in vRealize Hyperic are available for selection.

Write Custom Reports

This topic provides instructions for writing custom reports.

If you need reports that vRealize Hyperic does not provide, you can write your own. Report Center is based on JasperReports, and the templates are written in JasperReports' XML format, JRXML. To write a report template, vRealize Hyperic suggests you use the Java-based application iReport. Custom report templates must be in either JRXML format (generated from the content created in iReport) or compiled .jasper format. At this time compiled .jasper format is the only format supported when deploying new templates to vRealize Hyperic

To incorporate a new report template into Report Center

Prerequisites

Because vRealize Hyperic reporting functionality uses JasperReports and only supports reporting generation using SQL queries against the Hyperic database, writing custom report templates requires familiarity with both JasperReports and the Hyperic database schema. Assistance with the creation of custom Reports may be provided by our Professional Services Organization. Please contact your sales account representative.

Running SQL against the production Hyperic database is not recommended, as it may have an impact on performance or compromise database integrity. An inefficient or improperly configured report query could compromise the database or affect performance.

The best practice is to run reports in a test environment during development and against a replicated production database.

Procedure

- 1 Write a report template in JRXML or compiled Jasper format, being sure to specify any user-input parameters necessary for generating the report (for example, a date range).
- 2 Place the template (.jasper format only) in /hq-engine/hq-server/webapps/ROOT/WEB-INF/reportTemplates.
- 3 Log out of vRealize Hyperic and log back in.

The report appears in the report list.

Specifying Report Description and Report Parameters

Use the following descriptions and parameters for reporting.

In the JasperReport file for a report:

- Use a <property> element to define a description for the report. The text you specify will appear as the Report Description when you select a report on the Report Center page

- Use a <parameter> element to specify a report parameter. When you run the report, you will be queried to supply the parameter value. The following parameter types are supported:
 - java.util.Date
 - java.util.String

Views for Specific Resource Types

5

This section describes pages in the user interface that are specific to particular resource types.

These views are enabled by the product plugin that manage the resource type, and are present in the vRealize Hyperic user interface only if you have resources of that type under management.

- [GemFire View](#)

This section provides information about the various GemFire views available through vRealize Hyperic

- [HQ vSphere View](#)

The HQ vSphere page is an interface for monitoring and managing vSphere Hosts and VMs available in vRealize Hyperic.

- [tc Server Application Management View](#)

This page is present in the user interface only if you have the plugin.

- [Agent Control View](#)

This section discusses the various parts of the Agent Control view.

- [Spring Insight View](#)

This page describes the Spring Insight Applications page, available if you have Insight Operations under management.

GemFire View

This section provides information about the various GemFire views available through vRealize Hyperic

The GemFire plugin provides a live data user interface for viewing metrics in real-time. (As opposed to the Monitor tab for a GemFire component, which presents metrics that have been saved to the vRealize Hyperic database.)

For more information about GemFire, see the *vRealize Hyperic Resource Configuration and Metrics Guide*.

- [GemFire View for a Distributed System](#)

The GemFire View for a Distributed System displays the following information.

- [GemFire View for a Cache Server or Application Peer](#)

The GemFire View for a Cache Server or Application Peer displays server and region inventory properties and metrics.

- [GemFire View for a Gateway Hub](#)

The HQ GemFire View for a Gateway Hub displays inventory data and live measurements for the Gateway Hub, the Hubs and Regions it contains, and clients connected to the Hub.

GemFire View for a Distributed System

The GemFire View for a Distributed System displays the following information.

- Servers in the DS — The number of Cache Servers, Gateway Hubs and Application Peers in the DS.
- Gateways in the DS — If the DS is part of a multi-site deployment, the number of Gateways in the DS's Gateway Hub.
- Clients connected to the DS — The number of clients (Cache Servers, Gateways, or Application Peers) in other Distributed Systems that are connected to the DS.

The table in the middle of the GemFire View for a DS displays inventory properties and last reported metric values for each server (including Cache Servers, Gateway Hubs and Application Peers) the DS.

GemFire View for a Cache Server or Application Peer

The GemFire View for a Cache Server or Application Peer displays server and region inventory properties and metrics.

- Server metrics — The single row table presents last reported metric values for the Cache Server or Application Peer.
- Region metrics — The multi-row table presents inventory properties and the Entry Count metric for each region the server contains.

GemFire View for a Gateway Hub

The HQ GemFire View for a Gateway Hub displays inventory data and live measurements for the Gateway Hub, the Hubs and Regions it contains, and clients connected to the Hub.

HQ vSphere View

The HQ vSphere page is an interface for monitoring and managing vSphere Hosts and VMs available in vRealize Hyperic.

The Performance tab appears for a VM only if there is a vRealize Hyperic agent running in the VM.

Display the HQ vSphere Page

To display the HQ vSphere page, select HQ vSphere from the **Resources** tab in the Masthead.

Note The HQ vSphere option appears on the Resources tab only in vRealize Hyperic, and only if you have vSphere components under management.

You can also navigate to the HQ vSphere page from the Resource Hub. The resource page for a vCenter server or vSphere Host has a View in HQ vSphere link above the Monitor tab.

HQ vSphere and Resource Permissions

In vRealize Hyperic, a user can only access resources that are assigned to groups to which the user's role grants access. The accessible resources are further limited by the role's permissions to inventory types.

In other words, you can see PlatformA in browse dialogs and navigate to it if: (1) it is a member of a group assigned to your role, and (2) that role grants access to platforms.

The HQ vSphere user interface behaves somewhat differently: if you have view access to platforms, you can see all of the vSphere Hosts and vSphere VMs in the deployment whether or not they belong to a group assigned to your role.

Note also that even if your role does not grant view permission to servers, vCenter servers will appear in the HQ vSphere page. A vCenter server is the root of the vSphere resource hierarchy and is visible regardless of role permissions; the only resource data exposed for the vCenter server is its name.

- [Information Available in the Inventory Tab](#)

This topic provides information about the Inventory Tab.

- [Information Available in the Summary Tab](#)

The **Summary** tab displays properties for the selected resource, and its parent resource, as applicable.

- [Information Available in the Performance Tab](#)

This topic provides information about what is available in the Performance tab.

- [Information Available in the Control Tab](#)

This topic provides information on the data available in the Control tab.

- [Metrics](#)

This topic lists the metrics for VMware vSphere hosts.

Information Available in the Inventory Tab

This topic provides information about the Inventory Tab.

The **Inventory** tab on the HQ vSphere page is a tree of the vSphere resources under HQ management, organized in the resource type hierarchy shown below. The lowest level appears for a VM that has an Hyperic Agent running and monitoring resources running in the VM.

The contents are updated once per minute.

You can use the tab to view the virtual resource hierarchy and to navigate among resources. When you select a resource, the tab or tabs on the right side of the page contain resource data, performance charts, and resource control commands, as appropriate to the resource type. The table below shows the vSphere resource type hierarchy; the right column indicates the inventory level for a type in the Hyperic inventory model.

Information Available in the Summary Tab

The **Summary** tab displays properties for the selected resource, and its parent resource, as applicable.

The **Summary** tab is available when any resource in the **Inventory** tab is selected.

Summary Tab for vCenter

When a vCenter Server instance is selected in the vSphere page, the **Summary** tab contains the name of the vCenter instance. You can view inventory and configuration properties for the vCenter instance in the Resource Hub — click view resource next to the resource name to view it in the Resource Hub.

Summary Tab for vSphere Hosts

The Summary tab for a vSphere Host displays the following properties:

- Host Information
 - Hostname
 - Location
 - Manufacturer
 - Model
 - VMWare Version
- Processor Details
 - Type of processor
 - CPU - Sockets and cores per socket
- Network Details
 - IP address
 - Default gateway

- DNS - primary and secondary DNS server

Summary Tab for vSphere VM

Summary Tab for vSphere VM The Summary tab for a vSphere VM displays the following VM properties:

- VM Information
 - Hostname - of the vSphere Host (ESX platform) where the VM runs
 - Guest OS - operating system running in the VM
 - vCPU(s) - number of virtual processors in the VM
 - Memory - VM memory, in MB
 - MAC Address
 - IP - VM's IP address
 - VM Version - virtual machine hardware version
 - Tools Version - version of VMware Tools on the VM.
- Config Details
 - ESX Host - IP address of the vSphere Host (ESX platform) where the VM runs
 - Resource Pool - resource pool with which the VM is associated
 - Config File - path to the the VM configuration (.vmx) file, expressed using the symbolic link path to the VMFS volume where the file is stored

Properties for the vSphere Host where the VM runs, described above in the Summary Tab for vSphere Hosts, are shown below the VM properties.

If the VM does not have an Hyperic Agent running in it, the Summary tab has a "Performance data not available" near the top of the page.

Summary Tab for a Managed Resource in the VM

When you select a managed resource running in a VM, the **Summary** tab displays the vSphere Host and VM properties described in Summary tab for vSphere VM.

The contents of the Summary tab are updated once per minute.

To view inventory properties for the managed resource itself, click the **view resource** link next to the resource name to view the resource in the Resource Hub.

Note This information applies to a VM with a running Hyperic Agent that is managing resources in the VM. Otherwise, resources running in the VM do not appear in the HQ vSphere page.

Information Available in the Performance Tab

This topic provides information about what is available in the Performance tab.

The **Performance** tab appears when a vSphere Host is selected, and, if the VMs running on the host have vRealize Hyperic agents running, for each VM, and for the managed resources running in the VMs.

The contents of the Performance tab are updated once per minute.

If a VM does not have an agent running in it, no Performance tab appears. You can view the VM metrics in its Monitor page in the Resource Hub. Click **view resource** next to the **Hostname** property on the VM's Summary tab.

View Metrics

The Performance tab displays an Availability bar, and a chart for each metric currently enabled for the selected resource type.

The tables in VMware vSphere Host Metrics list supported vSphere metrics and default settings.

Twelve-hour history is displayed by default. You can use the Data Range pull-down to set the display range to the most recent:

- 1, 4, or 12 hours
- 1 or 2 days,
- 1 week, or
- 1 month

Correlate Metrics

On the **Performance** tab for a VM or an vRealize Hyperic agent-managed resource running in the VM, you can use the **Compare** pull-down to correlate the selected resource's performance with its parent or grandparent. You can:

- Compare each VM metric with its vSphere Host equivalent — For example compare the VM's "Disk Usage (Average)" metric to its host's "Disk Usage (Average)" metric; the VM's "CPU Usage (Average)" metric to its host's "CPU Usage (Average)" metric; and so on.
- Compare each VM metric to a selected vSphere Host metric — For example, compare each VM metric to the "Disk Usage (Average)" metric of its vSphere host.
- Compare metrics for a managed resource in the VM with performance of the VM or the vSphere host — Use the Compare pull-down to select a VM or Host metric - the charts for the managed resource metrics will be overlaid with the line for the selected VM or host metric.

Information Available in the Control Tab

This topic provides information on the data available in the Control tab.

The **Control** tab appears when a VM is selected (if there is an agent running in the VM) and allows you to issue a control command to the VM.

If a VM does not have an agent running in it, no **Control** tab appears. You can run control actions from the Control page for the VM in the Resource Hub. Click **view resource** next to the Hostname property on the VM's Summary tab.

Metrics

This topic lists the metrics for VMware vSphere hosts.

Table 5-1. VMware vSphere Host Metrics

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Availability	Availability	n/a, Availability is determined by HQ; it is not a vCenter statistic.	<p>The plugin determines a vSphere Host's availability by querying vCenter for the host's power state, once per minute, by default. Availability takes these values: If the power state is:</p> <ul style="list-style-type: none"> ■ Up - if vSphere Host power state is powered On. ■ Paused - if vSphere Host power state is standby. ■ Powered Off - if vSphere Host power state is powered Off. ■ Down - if the vSphere Host power 	Percent	Availability	True	1 min

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
			state is none of the above.				
Uptime	sys.uptime.latest			sec	Availability	False	1 min
CPU Usage (Average)	cpu.usage.average	1	CPU usage as a percentage during the interval. Actively used CPU of the host, as a percentage of the total available CPU. Active CPU is approximately equal to the ratio of the used CPU to the available CPU. available CPU = # of physical CPUs × clock rate 100% represents all CPUs on the host. For example, if a four-CPU host is running a virtual machine with two CPUs, and the usage is 50%, the host is using two CPUs completely.	Percent	Utilization	True	5 min

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Usage (Average)	mem.usage.average	1	memory usage = memory consumed ÷ host configured memory size	Percent	Utilization	True	5 min
Network Usage (Average)	net.usage.average	1	Sum of the data transmitted and received during the collection interval. Sum of data transmitted and received across all physical NIC instances connected to the host.	KB	Utilization	True	5 min
CPU Reserved Capacity	cpu.reservedCapacity.average	2	Total CPU capacity reserved by the virtual machines.	none	Utilization	False	5 min.
CPU Throttled (1 min. Average)	rescpu.maxLimited1.latest	3		percent	Utilization	False	5 min.
CPU Throttled (5 min. Average)	rescpu.maxLimited5.latest	3		percent	Utilization	False	5 min.
CPU Throttled (15 min. Average)	rescpu.maxLimited15.latest	3		percent	Utilization	False	5 min.
CPU Running (1 min. Average)	rescpu.running1.latest	3		percent	Utilization	False	5 min.
CPU Running (5 min. Average)	rescpu.running5.latest	3		percent	Utilization	False	5 min.
CPU Running (15 min. Average)	rescpu.running15.latest	3		percent	Utilization	False	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
CPU Active (1 min. Average)	rescpu.actav 1.latest	3		percent	Utilization	False	5 min.
CPU Active (5 min. Average)	rescpu.actav 5.latest	3		percent	Utilization	False	5 min.
CPU Active (15 min. Average)	rescpu.actav 15.latest	3		percent	Utilization	False	5 min.
Memory Swap In	mem.swapin. average	2	Sum of memory swapin of all powered on VMs on the host.	KB	Utilization	False	5 min.
Memory Swap Out	mem.swapout. average	2	Sum of Memory Swap Out of all powered on VMs on the host.	KB	Utilization	False	5 min.
Memory Swap Used	mem.swapused. average	2	Amount of memory that is used by swap. Sum of Memory Swapped of all powered on virtual machines and vSphere services on the host.	KB	Utilization	False	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Balloon	mem.vmmem ctl.average	1	Sum of Memory Balloon (Amount of memory allocated by the virtual machine memory control driver) of all powered on virtual machines and vSphere services on the host. If the balloon target value is greater than the balloon value, the VMkernel inflates the balloon, causing more virtual machine memory to be reclaimed. If the balloon target value is less than the balloon value, the VMkernel deflate the balloon, which allows the virtual machine to consume additional memory if needed.	KB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Unreserved	mem.unreserved.average	2	Amount of memory that is unreserved. Memory reservation not used by the Service Console, VMkernel, vSphere services and other powered on VMs' user-specified memory reservations and overhead memory.	KB	Utilization	True	5 min.
Memory Heap	mem.heap.average	2	Amount of VMkernel virtual address space dedicated to VMkernel main heap and related data.	KB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Heap Free	mem.heapfree.average	2	Amount of free address space in the VMkernel's main heap. Heap Free varies, depending on the number of physical devices and various configuration options. There is no direct way for the user to increase or decrease this statistic.	KB	Utilization	True	5 min.
Memory Overhead	mem.overhead.average	1	Total of all overhead metrics (Amount of additional machine memory allocated to a virtual machine for overhead. The overhead amount is beyond the reserved amount).for powered-on virtual machines, plus the overhead of running vSphere services on the host.	KB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Zero	mem.zero.average	2	Amount of memory that is zeroed out (contains only 0s). This statistic is included in Memory Shared. For a vSphere Host, Sum of Memory Zero of all powered on VMs and vSphere services on the host.	KB	Utilization	True	5 min.
Memory Reserved Capacity	mem.reserve.dCapacity.average	2	Total amount of memory reservation used by powered on VMs and vSphere services on the host. Includes overhead amount.	MB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Active	mem.active.average	2	Amount of memory actively used, as estimated by VMkernel. Active memory is based on the current workload of the virtual machine or host. For a vSphere Host, sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMKernel applications on the host.	KB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Shared	mem.shared.average	2	Sum of the shared memory values of all powered-on virtual machines, plus the amount for the vSphere services on the host. The host's Memory Shared may be larger than the amount of machine memory if memory is overcommitted (the aggregate virtual machine configured memory is much greater than machine memory). The value of this statistic reflects how effective transparent page sharing and memory overcommitment are for saving machine memory.	KB	Utilization	True	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Granted	mem.granted.average	2	The total of all granted metrics for all powered-on virtual machines, plus memory for vSphere services on the host.	KB	Utilization	False	5 min.
Memory Consumed	mem.consumed.average	1	Amount of machine memory used on the host. Consumed memory includes memory used by virtual machines, the service console, VMkernel, and vSphere services, plus the total consumed memory for all running virtual machines. host consumed memory = total host memory - free host memory	KB	Utilization	False	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory State	mem.state.lat est	2	<p>Amount of free machine memory on the host. VMkernel has four free-memory thresholds that affect the mechanisms used for memory reclamation.</p> <ul style="list-style-type: none"> ■ 0 (High) - Free memory \geq 6% of machine memory - service console memory ■ 1 (Soft) - Free memory \geq 4% of machine memory - service console memory ■ 2 (Hard) - Free memory \geq 2% of machine memory - service console memory ■ 3 (Low)- Free memory \geq 1% of machine 	None	Utilization	False	5 min.

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
			memory - service console memory				
			For 0 and 1, swapping is favored over ballooning.				
			For 2 and 3, ballooning is favored over swapping.				

Table 5-1. VMware vSphere Host Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Shared Common	mem.sharedcommon.average	2	Amount of machine memory that is shared by all powered-on virtual machines and vSphere services on the host. Memory Shared - Memory Shared Common = Host memory saved by sharing	KB	Utilization	False	5 min.
Memory Used by vmkernel	mem.sysUsage.average	2	Amount of memory used by the VMkernel. Amount of machine memory used by the VMkernel for "core" functionality (such as its own internal uses, device drivers, etc). It does not include memory used by VMs or by vSphere services.	KB	Utilization	False	5 min.

VMware vSphere VM Metrics

Table 5-2. VMware vSphere VM Metrics

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Availability	Availability	n/a, Availability is determined by HQ; it is not a vCenter statistic.	<p>The plugin determines a vSphere Host's availability by querying vCenter for a VM power state, once per minute, by default. Availability takes these values: If the power state is:</p> <ul style="list-style-type: none"> ■ Up - if VM power state is powered On. ■ Paused - if VM power state is standby. ■ Powered Off - if VM power state is powered Off. ■ Down - if the VM power state is none of the above. 	Percent	Availability	True	1 min
Uptime	sys.uptime.la test	1		sec	Availability	False	1 min

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
CPU Usage (Average)	cpu.usage.average	1	CPU usage as a percentage during the interval. This value is reported with 100% representing all processor cores on the system. As an example, a 2-way VM using 50% of a four-core system is completely using two cores.	Percent	Utilization	True	5 min
Disk Usage (Average)	disk.usage.average	1	Aggregated disk I/O rate.	KB per second	Utilization	True	5 min
Memory Usage (Average)	mem.usage.average	1	The percentage of memory used as a percent of all available machine memory.	Percent	Utilization	True	5 min
Network Usage (Average)	net.usage.average	1	Sum of the data transmitted and received during the collection interval.	KB	Utilization	True	5 min.
CPU Throttled (1 min. Average)	rescpu.maxLimited1.latest	3		percent	Utilization	False	5 min.
CPU Throttled (5 min. Average)	rescpu.maxLimited5.latest	3		percent	Utilization	False	5 min.
CPU Throttled (15 min. Average)	rescpu.maxLimited15.latest	3		percent	Utilization	False	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
CPU Running (1 min. Average)	rescpu.runav1.latest	3		percent	Utilization	False	5 min.
CPU Running (5 min. Average)	rescpu.runav5.latest	3		percent	Utilization	False	5 min.
CPU Running (15 min. Average)	rescpu.runav15.latest	3		percent	Utilization	False	5 min.
CPU Active (1 min. Average)	rescpu.actav1.latest	3		percent	Utilization	False	5 min.
CPU Active (5 min. Average)	rescpu.actav5.latest	3		percent	Utilization	False	5 min.
CPU Active (15 min. Average)	rescpu.actav15.latest	3		percent	Utilization	False	5 min.
Memory Swap In	mem.swapin.average	2	Average memory Swap In. A large number here represents a problem with lack of memory and a clear indication that performance is suffering as a result.	KB	Utilization	False	5 min.
Memory Swap Out	mem.swapout.average	2	Average Memory Swap Out. A large number here represents a problem with lack of memory and a clear indication that performance is suffering as a result.	KB	Utilization	False	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Balloon	mem.vmmem ctl.average	1	Sum of Memory Balloon (Amount of memory allocated by the virtual machine memory control driver). If the balloon target value is greater than the balloon value, the VMkernel inflates the balloon, causing more virtual machine memory to be reclaimed. If the balloon target value is less than the balloon value, the VMkernel deflate the balloon, which allows the virtual machine to consume additional memory if needed.	KB	Utilization	True	5 min.
Memory Balloon Target	mem.vmmem ctltarget.average	2		KB	Utilization	False	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Zero	mem.zero.average	2	Amount of memory that is zeroed out (contains only 0s). This statistic is included in Memory Shared. For a vSphere Host, Sum of Memory Zero of all powered on VMs and vSphere services on the host.	KB	Utilization	False	5 min.
Memory Active	mem.active.average	2	Amount of memory actively used, as estimated by VMkernel. Active memory is based on the current workload of the virtual machine or host. For a vSphere Host, sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMKernel applications on the host.	KB	Utilization	True	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Shared	mem.shared.average	2	The average amount of shared memory. Shared memory represents the entire pool of memory from which sharing savings are possible. The amount of memory that this has been condensed to is reported in shared common memory. So, total saving due to memory sharing equals shared memory minus shared common memory.	KB	Utilization	True	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Granted	mem.granted.average	2	The amount of memory that was granted to the VM by the host. Memory is not granted to the host until it is touched one time and granted memory may be swapped out or ballooned away if the VMkernel needs the memory.	KB	Utilization	False	5 min.
Memory Consumed	mem.consumed.average	1	The amount of machine memory that is in use by the VM. While a VM may have been configured to use 4 GB of RAM, as an example, it might have only touched half of that. Of the 2 GB left, half of that might be saved from memory sharing. That would result in 1 GB of consumed memory.	KB	Utilization	False	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Shared	mem.shared.average	2	Sum of the shared memory values of all powered-on virtual machines, plus the amount for the vSphere services on the host. The host's Memory Shared may be larger than the amount of machine memory if memory is overcommitted (the aggregate virtual machine configured memory is much greater than machine memory). The value of this statistic reflects how effective transparent page sharing and memory overcommitment are for saving machine memory.	KB	Utilization	True	5 min.

Table 5-2. VMware vSphere VM Metrics (Continued)

Name	Alias	vCenter Statistics Level	Definition	Units	Category	Default On	Default Interval
Memory Granted	mem.granted.average	2	The total of all granted metrics for all powered-on virtual machines, plus memory for vSphere services on the host.	KB	Utilization	False	5 min.
Memory Consumed	mem.consumed.average	1	Amount of machine memory used on the host. Consumed memory includes memory used by virtual machines, the service console, VMkernel, and vSphere services, plus the total consumed memory for all running virtual machines. host consumed memory = total host memory - free host memory	KB	Utilization	False	5 min.

tc Server Application Management View

This page is present in the user interface only if you have the plugin.

- [tc Runtime Application Management](#)

The Application Management page enables an authorized user to manage applications deployed that server instance or group.

- [Navigate to a tc Runtime Instance Or Group](#)
Follow this process to access the Application Management page.
- [Create tc Runtime Groups](#)
Follow these instructions to create tc runtime groups.
- [Listing Applications](#)
This topic provides information about the list of deployed applications.
- [Deploy Applications](#)
Follow this process to deploy applications.
- [Undeploying Applications](#)
Follow this process to undeploy applications.
- [Start an Application](#)
This topic provides instructions for starting an application.
- [Stop an Application](#)
This topic provides instructions for stopping an application.
- [Reload an Application](#)
This topic provides instructions for reloading an application.

tc Runtime Application Management

The Application Management page enables an authorized user to manage applications deployed that server instance or group.

The Application Management page is available when a tc Runtime instance or a group of tc Runtime instances is selected.

The HQ super-user can use all the tc Runtime Application Management features of the user interface. If, however, you log on to the user interface as a non-super-user, then that user must have certain permissions to be able to use the tc Runtime features.

An HQ group is not equivalent or similar to a tc Runtime cluster.

- An HQ group is a set of other managed resources. An authorized user can create a group and assign resources to it. A group whose members are all of the same resource type is referred to as a compatible group. You can manage and monitor the resources in a compatible group at the group level. Grouping tc Runtime instances allows you manage applications across multiple tc Runtime instances.
- A tc Runtime cluster enables session replication. Clustering behavior is defined in a tc Runtime instance's server.xml file.

Navigate to a tc Runtime Instance Or Group

Follow this process to access the Application Management page.

To navigate to the Application Management page of a tc Runtime instance or group.

Procedure

- 1 Click **Resource > Browse** at the top of the HQ Dashboard.
 - a To browse to a server instance, click **Servers** to view a list of the servers to which you have access. tc Runtime instances have the server type "SpringSource tc Runtime 6.0". Apache Tomcat server instances have the server type Tomcat X.X.
 - b To browse to a group of servers, click **Compatible Groups/Clusters** to view a list all the compatible groups to which you have access. Groups of tc Runtime have the type SpringSource tc Runtime 6.0.
- 2 In the table, click the name of the tc Runtime instance or group to which you want to navigate.
- 3 Click **Views > Application Management**.

Create tc Runtime Groups

Follow these instructions to create tc runtime groups.

Grouping tc Runtime instances eases the process of managing server instances and applications. For example, you can deploy an application to group of tc Runtime with a single command, instead of deploying to each instance individually. To create a group of tc Runtime instances:

Procedure

- 1 Click **Resources > Browse** at the top of the dashboard.
- 2 Click **Servers** to list servers in inventory to which you have access.

Only server instances that have been auto-discovered and added to inventory appear. tc Runtime instances have the server type SpringSource tc Runtime 6.0. Apache Tomcat server instances have the server type Tomcat X.X.
- 3 Select the checkbox to the left of each tc Runtime instance you want to include in the new group.
- 4 Click **Group**.
- 5 Enter a name for the group (required) and a description and location as desired.
- 6 Click **OK**.

Listing Applications

This topic provides information about the list of deployed applications.

Deployed applications are listed in the Deployed Applications section of the the Application Management page for a server instance or group. For a group, the table lists all applications deployed to all servers in the group. The table shows:

- **Status** — The state of the application - either "Running" or "Stopped". For a group of servers:
 - "Running" or "Stopped" indicates that all instances of the application on all servers in the group have that state.
 - A numeric value indicates the number of servers upon which the application's state is "Running".
- **Sessions** — The number of current active sessions for the application.
- **Revision** — Current revision of the application. Hyperic creates a new revision each time you deploy an application that uses an existing context path

Note Only applications on running tc Runtime instances appear. HQ returns an error for application hosts that are not running.

Deploy Applications

Follow this process to deploy applications.

Deploying an application is the process of uploading it to tc Runtime and making it available to users. You can deploy an application to a single tc Runtime instance or to a group of tc Runtime instances. If you deploy to a group, the application will be deployed to each of tc Runtime instance in that group, enabling users to run the application from any of the tc Runtime instances in the group.

You can deploy an application from a WAR file located:

- on your local file system (on the the same machine as the browser you are using to connect to HQ's web user interface), or
- on the computer on which the tc Runtime resource is running.

Whether the WAR file is local or remote, you can either enter a new context path (string that uniquely identifies the Web application in the URL used to invoke it), or use the default value, which is the name of the WAR file, without the .war extension.

Procedure

- 1 Navigate to the Application Management page for a tc Runtime instance or group.
- 2 Depending on the location of the WAR file for the application, either:
 - Click Browse in the Deploy Application From Local Machine section to browse to the file's location.

- Enter the full path to the WAR file on the computer hosting the tc Runtime resource in the Deploy Application from Server Machine section.
- 3 As desired, enter a context path in the **Context path** text box.
 - 4 Optionally, select the **Use cold deployment strategy** checkbox if you want the tc Runtime instance to shutdown, deploy the application, and then start up again.

By default (if box is clear), the tc Runtime instance hot-deploys the application, which means it does not shutdown then restart but simply deploys the application while the instance is still running. Use the cold deployment strategy if you want to avoid common hot deployment errors, such as running out of PermGen space. The PermGen space holds the metadata about classes that have been loaded/created in the JVM.
 - 5 Click **Upload and Deploy** or **Deploy**, whichever is appropriate.

See the top Results of the last operation section for details about the result of deploying the application. When you deploy an application to tc Runtime, the application is started.

Undeploying Applications

Follow this process to undeploy applications.

Undeploying an application removes it from the tc Runtime instance or group. If you want to temporarily prevent users from access an application, stop it rather than undeploy it. For more information, see [Stop an Application](#).

Procedure

- 1 Navigate to the application management page of a tc Runtime instance or group.
- 2 In the Deployed Applications section, select the checkbox to the far-left of the application(s) you want to undeploy.
- 3 Click **Undeploy**. HQ removes the application from the list of deployed applications.

See the top Results of the last operation section for details about the results of undeploying the application.

Start an Application

This topic provides instructions for starting an application.

Prerequisites

You can only start applications that have been deployed. For more information, [Deploy Applications](#)

Procedure

- 1 Navigate to the application management page of a tc Runtime instance or group.
- 2 In the Deployed Applications section, select the checkbox to the far-left of the application(s) you want to start.

3 Click **Start**.

The status of the application changes to Running.

See the top Results of the last operation section for details about the result of starting the application.

Stop an Application

This topic provides instructions for stopping an application.

Stopping an application makes it unavailable to users.

Procedure

- 1 Navigate to the application management page of a tc Runtime instance or group.
- 2 In the Deployed Applications section, select the checkbox to the far-left of the application(s) you want to stop.
- 3 Click **Stop**.

The status of the application changes to Stopped.

See the top Results of the last operation section for details about the result of starting the application.

Reload an Application

This topic provides instructions for reloading an application.

When you reload an existing application, it shuts itself down and then reloads itself.

Procedure

- 1 Navigate to the application management page of a tc Runtime instance or group.
- 2 In the Deployed Applications section, select the checkbox to the far-left of the application(s) you want to reload.
- 3 Click **Reload**. The status of the application changes to Running.

See the top Results of the last operation section for details about the result of reloading the application.

Agent Control View

This section discusses the various parts of the Agent Control view.

You can run vRealize HypericAgent control commands — for a single agent, or for a group of agents — from the vRealize Hyperic user interface.

- [Restart an Agent or a Group of Agents](#)

This section provides instructions for restartnig an agent or group of agents.

- [Ping an Agent or a Group of Agents](#)

This topic provides instructions for checking if an agent or group of agents is responsive.

- [Upgrade an Agent or Group of Agents](#)

This topic provides instructions for upgrading an agent or group of agents.

- [Push a Plugin to an Agent or Group of Agents](#)

This topic provides instructions on pushing a plugin to an agent or group of agents.

Restart an Agent or a Group of Agents

This section provides instructions for restarting an agent or group of agents.

Procedure

- 1 Navigate to the vRealize Hyperic agent or group of agents.
- 2 Click the **Views** tab.
- 3 Click **Agent Commands**.
- 4 Select the restart command.
- 5 Click **Execute**.

This command invokes the agent's Java Service Wrapper's restart command, which shuts down the JVM process in which the agent runs, waits for the process to terminate cleanly, and spawns a new JVM process for the agent. During the restart process, the agent's metric collection and resource control functionality will be interrupted.

The restart command happens asynchronously. To verify successful restart, navigate to the agent in the vRealize Hyperic user interface check its availability.

Ping an Agent or a Group of Agents

This topic provides instructions for checking if an agent or group of agents is responsive.

Procedure

- 1 Navigate to the vRealize Hyperic agent or group of agents.
- 2 Click the **Views** tab.
- 3 Click **Agent Commands**.
- 4 Select the ping command.
- 5 Click **Execute**.

This command invokes the agent's Java Service Wrapper's ping command on the agent.

If the ping is successful, this message appears in the Result pane on the right side of the page:

```
Successfully sent ping command to agent with id nnnnn.
```

If the ping is not successful, this message appears in the Result pane on the right side of the page:

```
Failed to send ping command to agent with id nnnnn. Reason: Unable to connect to AgentIpAddress.
```

Upgrade an Agent or Group of Agents

This topic provides instructions for upgrading an agent or group of agents.

Procedure

- 1 Navigate to the vRealize Hyperic agent or group of agents.
- 2 Click the **Views** tab.
- 3 Click **Agent Commands**.
- 4 Select the upgrade command.
- 5 Click **Execute**.
 - The selected agent bundle is transferred from the Hyperic Server to the target agent(s).
 - The agent expands the bundle locally.
 - The agent updates the local bundle property.
 - The server restarts the agent.
 - The configuration settings in the `agent.properties` file are preserved.

Push a Plugin to an Agent or Group of Agents

This topic provides instructions on pushing a plugin to an agent or group of agents.

This operation allows you to transfer new, custom, updated, or patched plugins from the vRealize Hyperic server's `ServerHome/hq-engine/hq-server/webapps/ROOT/WEB-INF/hq-plugins` directory to the target agents' `AgentHome/bundles/AgentBuildDir/pdk/plugins` directory. Pushing a plugin to an agent results in an agent restart.

Procedure

- 1 Navigate to the vRealize Hyperic agent or group of agents.
- 2 Click the **Views** tab.
- 3 Click **Agent Commands**.
- 4 Select the push plugin command.
- 5 Select a plugin from the plugin pull-down list.
- 6 Click **Execute**.

Spring Insight View

This page describes the Spring Insight Applications page, available if you have Insight Operations under management.

Understanding Application Metrics

The Spring Insight Applications page presents multiple metrics that indicate how well an application is running, for instance, the percentage of application requests had an unacceptable response time, how many requests return errors, and so on.

These metrics, described in the sections below, appear in several places in the Spring Insight Applications user interface: on the Health tab, which presents metrics for all applications reporting to vRealize Hyperic monitored Dashboards applications, and also on the Performance tab for an Insight Dashboard or an Insight Application.

Application Health

The Application Health metric for an application is the health value that Insight calculates for the application. The Insight health metric is based on the percentage of Insight traces that complete without error. The values, indicators, and thresholds for the Application Health metric are defined below.

Application Vitals

The Vitals metric for an application is calculated by vRealize Hyperic. It is based on the Application Health metric (described above), and also takes into account the vRealize Hyperic metrics and events for the application and for resources that the application depends upon, including:

- Number of unfixed alerts for the application in vRealize Hyperic.
- Availability of the managed resources that the application depends upon: application servers, database servers (if applicable), and platforms.
- Number of unfixed alerts for the resources that the application depends upon:
 - application servers
 - database servers
 - platforms

If an application has no unfixed alerts in vRealize Hyperic and the resources it depends upon are available and have no unfixed alerts, the application's value of the applications Vitals metric is equal to the value of its Application Health metric.

Unfixed Alert Count

The number of unfixed alerts during the timeframe.

Error Rate

The percentage of Insight traces that resulted in an HTTP error status code (500 to 600).

Response Time Levels

The "Invocations" metric tells you what percentage of traces had response times in each of three response time levels. The response time thresholds for each level are configurable in Spring Insight, and by default are:

- Satisfied - Response time 200 ms or less.
- Tolerated - Response time over 200 ms but less than 800 ms.
- Frustrated - Response time 800 ms or over.

For example, 19.8%/7.4%/72.7% indicates that over the selected interval:

- 19.2% of the traces had response time 200 ms or less.
- 7.5% had response time over 200 ms but less than 800 ms.
- 73.4% had response time 800 ms or over.

Working with the Spring Insight Applications Page

This topic provides information about the different sections of the Spring Insight Application page.

Dashboard Tab

The Dashboards pane on the left side of the Spring Insight Applications page is a tree control whose nodes are Insight Dashboards and Insight Applications.

The color of the bar to the left of a Dashboard node indicates the worst Vitals metric among the applications reporting to the Dashboard.

You use the Dashboards pane to browse a hierarchy of performance data. Depending on your need and interest, you view filter your view to include health information for:

- All applications reporting to all Hyperic-managed Insight Dashboards
 - This is the view that appears when neither a dashboard nor application is selected in the tree control. To return to this view, click the **Dashboards** link above the tree control.
- Just the applications reporting to a particular Insight Dashboards
 - To display this view, click the dashboard of interest in the tree control.
- Just one particular application
 - To display this view, select the application of interest in the tree control.

Health Tab for all Applications

The **Health** tab summarizes the health of all of the applications reporting to vRealize Hyperic-managed Insight Dashboards.

This tab is present when you first open the Spring Insight Applications page. After navigating to a dashboard or an application, you can redisplay the **Health** tab by clicking the **Dashboards** link at the top of the left pane. The **Health** tab contains the following data and metrics:

- Applications Health Trend chart — By default, this chart plots the health of the five applications that currently have the poorest vitals. You can add and remove applications to the chart by checking and unchecking the box next to it.
- Worst Vitals — The Vitals indicator in the All Applications bar reflects the worst value reported among all applications reporting to all Dashboards.
- Application Metrics — Below the Applications Health Trend chart is a list of all applications under management, ordered by the applications' Vitals metric, starting with the worst. The following data and metrics are displayed for each application:
 - Application — The name of the application is a link; if you click it, the **Performance** tab for the dashboard appears.
 - Vitals
 - Alerts
 - Errors
 - Invocations (F/T/S)

Performance Tab for a Dashboard

To display the **Performance** tab for an Insight Dashboard, click the dashboard's name in the **Dashboards** tab.

In the list of applications reporting to the dashboard:

- The list is ordered by application health, from worst to best.
- You can mouse over the application name to highlight the line for it in the Application Health Trend chart.
- You can click the application name to display the Performance tab for it.

The **Performance** tab for a Dashboard contains:

- Applications Health Trend chart — By default, this chart plots the health of the five applications (reporting to the Dashboard) that currently have the poorest vitals. You can add and remove applications to the chart by checking and unchecking the box next to it.
- Dashboard Vitals — The Vitals indicator in the Dashboard bar is a composite health indicator that takes into account:
 - The Vitals for the applications reporting to the currently selected Insight Dashboard.
 - Availability of the managed resources that the Insight Dashboard depends upon: application server (tc Server) and platform.
 - Number of unfixd alerts for the resources that the Insight Dashboard depends upon: application server (tc Server) and platform.

- Application Metrics — Below the Applications Health Trend chart is a list of the applications reporting to the Dashboard. The following metrics are displayed for each application:
 - Vitals
 - Alerts
 - Errors
 - Invocations (F/T/S)
- Application Server Performance — To view health data for the application server that the Insight Dashboard runs on, click the control in the Application Server bar.
- Platform Performance — To view health data for platform on which the Insight Dashboard runs, click the control in the Platforms bar.

Application Server Performance

The Application Server section charts a selected app server metric for all, or selected application servers where the applications in the dashboard run.

The **Metric Low/Avg/Peak** pulldown lists all of the available metrics for the application server and shows the low, average, and peak values for the metric over the current time range. Choose a metric from the list to graph it.

Below the chart, the following information is listed for each application server:

- name - The name of the resource in vRealize Hyperic.
- AVAIL - Current availability of the resource.
- ALERTS - Current number of unfixed alerts.
- LOW - Low value during current display range.
- AVG - Average value during current display range.
- PEAK - High value during current display range.

Platform Performance

The Platform section charts a selected platform metric for all, or selected platforms where the applications in the dashboard run.

The **Metric Low/Avg/Peak** pulldown lists all of the available metrics for the application server and shows the low, average, and peak values for the metric over the current time range. Choose a metric from the list to graph it.

Below the chart, the following information is listed for each application server:

- name - The name of the resource in vRealize Hyperic.
- AVAIL - Current availability of the resource.
- ALERTS - Current number of unfixed alerts.

- LOW - Low value during current display range.
- AVG - Average value during current display range.
- PEAK - High value during current display range.

Performance Tab for an Application

To display the Performance tab for an Insight Application, click the application's name in the Dashboards tab.

The Performance tab has four sections:

- Applications
 - The application name is a link - if you click it, Insight Operations will open in a new window, and display more detailed resource health and performance metrics for the current display range.
 - The Vitals health bar to the right of the application name indicates the current Vitals for the application.
 - The Health Trend chart plots the application's "Health" and "Error Rate" metrics for the display range.
 - The Invocations chart plots the distribution of request response times across the three response time levels
- Application Server
 - This section is expanded if one or more application servers to which the app is deployed is unhealthy. Otherwise, the section is collapsed.
- Data Service
 - This section is present if the application access a Hyperic managed database server.
 - This section is expanded if the database that the app uses is unhealthy application. Otherwise, the section is collapsed.
- Platform
 - This section is expanded if one or more platforms where the app runs is unhealthy. Otherwise, the section is collapsed.

The Data Service section is present if the currently selected application accesses a database managed by vRealize Hyperic.

The **Metric Low/Avg/Peak** pulldown lists all of the available metrics for the application server and shows the low, average, and peak values for the metric over the current time range. Choose a metric from the list to graph it.

Below the chart, the following information is listed for each application server:

- name - The name of the resource in vRealize Hyperic.
- AVAIL - Current availability of the resource.

- ALERTS - Current number of unfixed alerts.
- LOW - Low value during current display range.
- AVG - Average value during current display range.
- PEAK - High value during current display range.

Running Live System and Metric Queries

6

This section provides information about running live system and metric queries.

- [Creating JMX MBean Queries for JVMs](#)

The JMX MBean Query tool allows you to search for MBeans, display their attributes, and invoke selected MBean operations on them.

- [Running Live Exec Queries](#)

vRealize Hyperic's Live Exec view allows you to run system commands on a managed platform, or a group of platforms.

Creating JMX MBean Queries for JVMs

The JMX MBean Query tool allows you to search for MBeans, display their attributes, and invoke selected MBean operations on them.

When vRealize Hyperic alerts you of availability or health issues with a JVM resource or service, you can use the JMX MBean Query tool to troubleshoot and resolve the problem.

The MBean attributes displayed in the JMX MBean Query tool are read-only; the console does not support attribute editing. In this version of vRealize Hyperic, support for MBean operations is limited to operations with one or no arguments of primitive or simple type.

The JMX MBean Query tool is available on the **Views** tab when you have a JVM selected in the HQ portal. These instructions assume that your JVM is in HQ inventory.

Note JVMs are not auto-discovered by vRealize Hyperic. You must add a JVM to vRealize Hyperic inventory and configure for monitoring as described in Sun JVM.

Create a Query

To access specific MBean attributes and operations, you enter search patterns in the respective text boxes

From the JMX MBean Query tool page:

Procedure

- 1 Enter values in the relevant fields:
 - Object Name Pattern - To specify the MBean or Mbeans of interest, enter a valid MBean object name pattern.
 - Attribute Regex Filter - To restrict the attributes returned, enter a valid Java regular expression.
 - Operation Regex Filter - To restrict the MBean operations returned, enter a valid Java regular expression.
- 2 Click the **Query MBeans** button to execute the search.

Enable Automatic Refresh

This topic provides instructions for defining the automatic refresh interval.

Procedure

- ◆ To periodically refresh the attribute values, select an interval from the **Refresh Interval** pulldown.

Create Saved MBean Searches

If you expect to perform the same search frequently, you can create an XML file that specifies the search criteria.

Saved Mbean searches use a file that you store it in a specific directory in your vRealize Hyperic installation. Each saved search will appear in the MBean Query page's **Preset Searches** menu.

When you run a saved search, the values you defined for the search **Object Name Pattern**, **Attribute Regex Filter**, and **Operation Regex Filter** appear in the text boxes at the top of the view.

Procedure

- 1 Create a saved search in accordance with the schema described below in [Schema for JMX Search Criteria], using the following convention:

SearchID-filter.xml

where *SearchID-filter* is a meaningful name for the search to which you append "-filter.xml" identifies the search.

The filters values defined in this file will then appear in the **Preset Searches** drop-down menu and will be named according to the filter ID.

- 2 Save the file in the following directory:
 ServerHome/hq-engine/hq-server/webapps/ROOT/hqu/jmx/conf

Schema for JMX Search Criteria

This topic provides information about the makeup of a JMX search criteria schema.

A saved MBean search criteria consists of a top-level filters element, containing at least one filter sub-element. Each filter sub-element consists of:

- An *id* attribute containing an identifier for the search. This string will appear in the **Preset Searches** drop-down menu. Each search you define and save on a vRealize Hyperic server must have a unique id attribute.
- An *objectName* sub-element, whose value is the object name pattern used for the MBean query.
- An *attributeRegex* subelement, whose value is the Java regular expression used to filter the attributes returned from the MBean query.
- An *operationRegex* subelement, whose value is the Java regular expression used to filter the operations returned from the MBean query.

For example:

```
<?xml version="1.0" encoding="UTF-8"?>

<filters>

  <filter id="JVM Memory MBean">
    <objectName>java.lang:type=Memory</objectName>
    <attributeRegex>.*Usage</attributeRegex>
    <operationRegex>.*</operationRegex>
  </filter>

  <filter id="JVM Platform MBeans">
    <objectName>java.lang:*</objectName>
    <attributeRegex>.*</attributeRegex>
    <operationRegex>.*</operationRegex>
  </filter>

</filters>
```

Running Live Exec Queries

vRealize Hyperic's Live Exec view allows you to run system commands on a managed platform, or a group of platforms.

Live Exec uses vRealize Hyperic SIGAR — an API for gathering system information in real-time. The table below describes the queries you can run from Live Exec.

Table 6-1. Live Exec Queries

Command	Data Returned
cpuinfo	All CPUs on a platform.
cpuperf	Usage percentages on each CPU and related data. Equivalent to the uptime command.
df	All filesystems on a platform. This does not list disk-usage percentages.
ifconfig	Network statistics on each of the platform's network interfaces.
netstat	Active socket connections on the platform.
top	All processes (that the vRealize Hyperic agent can see) running on the platform. If a process is not listed, it is most likely because the agent does not have permission to see it.
who	Logged-in users

Run Live Exec on a Platform or Group of Platforms

This topic provides instructions for running Live Exec.

Procedure

- 1 Navigate to a platform, or a group of platforms in the Resource Hub.
- 2 Click the resource's **Views** tab.
- 3 Click **Live Exec**.

Available queries are listed in the lefthand page of the Live Exec page.

If the selected resource is a group, the individual platforms in the group appear as well. You can select individual platforms within the group.

- 4 Select a command from the **Please select a query to run** pulldown.

Command results appear in the righthand side of the page.

- If the resource is an individual platform, the results pane includes its name, with the name of the platform and the selected command listed about the data table.
- If the resource is a group:
 - If a member of the group has not previously been selected, click next to a specific platform now. That platform will be displayed first by default for subsequent commands.
 - If a member of a group is not available, click next to the resource to display the reason why it is unavailable.
 - To view the real-time data for another platform in the group, click next to that platform. The platform for which data is currently being displayed is highlighted at the left.

If the agent cannot reach the platform, an error message is displayed.

Using Resource Control Actions

This section provides information about using resource and control actions.

- [Running Resource Control Actions](#)

This section has information for running control actions on resources that support control actions.

- [Viewing Control Action Status and History](#)

This topic has information about how to view control action history.

Running Resource Control Actions

This section has information for running control actions on resources that support control actions.

There are two types of control actions in vRealize Hyperic: control actions enabled by the managing resource plugin, and custom control actions implemented by vRealize Hyperic administrators.

The process you use to selecting a control action to run varies for built-in and custom control actions:

- You select built-in control actions from the **Control** tab of the target resource — generally a server type.
- A custom control action is configured by a vRealize Hyperic administrator as service (of type "FileServer Service") on the platform where it runs. So, even if a custom control action target a server running on the platform, to invoke it, you navigate the platform service that represents the service on the parent platform.

Run a Built-in Control Action

To run a built-in control action - one that is enabled by a resource plugin - you navigate to the target resource, usually some type of server. (Control actions are typically available for servers, not for platforms, and rarely for services.)

Procedure

- 1 Navigate to the target resource.
- 2 Navigate to the **Control** tab.
- 3 In the Quick Control section of the page, use the **Control Action** pull-down to display a list of supported actions.

- 4 Select the desired control action, and click the green arrow to the right to initiate it.

Run a Control Action from a List of Custom Controls

This topic describes how to run a control action from a list of custom controls.

Procedure

- 1 To select from a list of all custom control actions, click **Browse > Services > FileServer File**.

OR

To select from a list of control actions for the resources on a platform, navigate to the platform. Custom control actions for a resources are listed in the Platform Services Health section.

- 2 Click the desired control action.
The **Monitor** tab for the custom control action appears.
- 3 Click the **Control** tab.
- 4 In the Quick Control section of the page, select **Run** from the **Control Action** pull-down to run the control action.

Schedule a Control Action

You can schedule a control action to be run in the future - either once, or on a recurring basis.

Procedure

- 1 Select the **Control** tab:
 - To schedule control actions for a resource with built-in control commands, navigate to the target resource, as described in [Run a Built-in Control Action](#).
 - To schedule control actions for a resource with built-in control commands, navigate to the target resource, as described in [Run a Control Action from a List of Custom Controls](#).
- 2 In the Control Action Schedule section of the page, click **New**.
- 3 On the New Scheduled Control Action Page select a control action from the **Control Action** pulldown.
- 4 Enter values in the following fields:
 - Description
 - Start — Specify when you want to run the control action (or kick off a schedule of periodic control action execution):
 - Click **immediately** if you want vRealize Hyperic to initiate the control action as soon as you save the schedule.
 - Otherwise, select the date and a time you want vRealize Hyperic to run the control action.

The **Recur** pulldown appears.

- 5 Determine the frequency with which vRealize Hyperic repeats the control action:
 - Never - select this option if you only want to run the command once.
 - Daily
 - Weekly
 - Monthly
- 6 Click **OK**.

Viewing Control Action Status and History

This topic has information about how to view control action history.

The Control Actions portlet displays information about recent control actions performed on resources over a configurable period of recent history.

- **Recent Control Actions** - This section lists the individual resource control action performed during the configured period, starting with the most recent. The following information is displayed for each resource:
 - **Resource Name** - Name of the resource.
 - **Control Action** - The control action that was performed.
 - **Date/Time** - When the action was performed
 - **Message** - Results of the control action.
- **Quick Control Frequency** - This section lists the resources upon which the most on-demand control actions have been performed during the configured interval. The following information is displayed for each resource:
 - **Resource Name** - Name of the resource.
 - **# of Control Actions** - How many control actions were performed during the interval.
 - **Most Frequent Control Action** - The control action that was most frequently performed.

Configure the Control Actions Portlet

This topic provides instructions for determining the amount of data presented in the Control Actions portlet.

Procedure

- 1 Click the gear icon in the upper right corner of the Control Actions portlet.
The Dashboard Settings: Control Actions page appears.

- 2 Configure what appears in the Recent Control Actions section using the first row of filters:
 - a Select the **Control Action Range** checkbox if you want to list recent control actions in the portlet. (To stop the display of this data, clear the checkbox.)
 - b In the first drop-down list, select the maximum number of completed control actions to list.
 - c In the second drop-down list, select the time range from which to gather that data.
- 3 Configure what appears in the Quick Control Frequency section using the second row of filters:
 - a Select the checkbox if you want to list the resources with the most on-demand control actions. (To stop the display of on-demand control-action data, clear the checkbox.)
 - b Select the maximum number of resources to list.
- 4 Click **OK**.

View Control History for a Resource

This topic provides instructions for viewing a list of control actions that were performed on a resource.

Procedure

- 1 Navigate to the resource in the **Resource Hub**.
- 2 Click the **Control** tab, which is present only if vRealize Hyperic supports control action for the selected resource type.
- 3 Click **History**.

The Control History page lists the following information for actions that were performed on the resources:

- Control Action — The action performed.
- Arguments — Arguments supplied to the control action.
- Command State — Status of the action: Completed, In Progress, or Failed.
- Date Started — Date and time that the control action was initiated.
- Elapsed Time — How long it took to perform the control action.
- User — User name of the vRealize Hyperic user who initiated or scheduled the control action.
- Message — (optional) Additional information about the control action.

To delete a control action, place a checkmark next to it in the list, and click **Delete** at the bottom of the page.

View Alerts and Update Alert Status



This section describes options for displaying the status of fired alerts, marking an alert "fixed", and acknowledging an alert that has an escalation.

- [Listing Recent Alerts](#)

The Dashboard's Recent Alerts portlet displays alerts that have recently been fired. You can limit the alerts shown by affected resource, priority, and quantity.

- [List Alerts for a Resource](#)

This topic describes how to produce a list of alerts for a resource.

- [Viewing Alerts Across All Resources](#)

This topic describes where you can view fixed and unfixed alerts.

- [Monitoring Unfixed Alerts](#)

This topic discusses the different ways you can monitor an unfixed alert.

- [View Alert Details](#)

This topic describes how to access alert details.

- [Acknowledging or Fixing an Alert](#)

This topic describes where to acknowledge or fix alerts.

Listing Recent Alerts

The Dashboard's Recent Alerts portlet displays alerts that have recently been fired. You can limit the alerts shown by affected resource, priority, and quantity.

List Alerts for a Resource

This topic describes how to produce a list of alerts for a resource.

The Alerts List page for a resource lists key information and status about alerts that have fired for the resource. By default, alerts that have fired on the current day are listed. You can select a different day to view.

Procedure

- 1 Browse to the resource.
- 2 Select the **Alert** tab.

Viewing Alerts Across All Resources

This topic describes where you can view fixed and unfixed alerts.

vRealize HypericAlert Center, available on the Analyze tab in the Masthead, provides system-wide and filtered views of fixed and unfixed alerts. You can include all alerts or those fired during a specified period, and filter by:

- whether the alert applies to an individual resource or a resource group
- minimum priority
- whether the alert is not fixed or in escalation

Monitoring Unfixed Alerts

This topic discusses the different ways you can monitor an unfixed alert.

The Operations Center page, available on the Analyze tab in the Masthead, provides system-wide and filtered views of unfixed alerts. You can filter by platform or group, and whether or not the alerts have escalations

View Alert Details

This topic describes how to access alert details.

There are several pages from which you can access the Alert Detail page for a fired alert:

Procedure

- ◆ Click the alert name on the Alert List page. Alternatively, you can
Click the alert definition name on the Alert Center page.

Acknowledging or Fixing an Alert

This topic describes where to acknowledge or fix alerts.

There are several pages from which you can mark an alert as Fixed, or acknowledge an alert that has an escalation:

- Recent Alerts portlet on the Dashboard
- • Alert List page
- • Alert Detail page
- Alert Center page