

vRealize Hyperic Administration 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 © 2016 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About vRealize Hyperic Administration Guide	6
1 Discover and Import Resources to Inventory	7
Import or Skip Resources in Auto-Discovery Portlet	7
Using the Auto-Discovery Results Page	8
Import or Skip Resources in Auto-Discovery Results Page	9
2 Options for Running and Controlling Resource Discovery	10
Configure the Frequency of the Default Scan	10
Configure the Frequency of the Runtime Scan	11
Scan a Platform On-Demand	11
Solving Auto-Discovery Problems	12
3 Setting Up Alert Notifications and Escalations	13
Tailoring Alert Notification Templates	13
Configuring Roles for Role-Based Alert Notifications	16
Enabling SNMP Trap Notifications	16
Enable Syslog Notifications	19
Configuring and Managing Escalations	20
Setting Up Alert Action Scripts	24
Configuring a Custom Control Action	25
Environment Variables for Fired Alert Data	27
4 Configure Monitoring and Alerting for a Resource	29
Configure Metric Baselines	29
Setting Up Log Tracking for a Resource	31
Setting Up Configuration Tracking for a Resource	35
5 Configuring Monitoring and Alerting for a Resource Type	39
Tailoring Metric Collection for a Resource Type	39
Define a Resource Type Alert	41
Define a Recovery Alert for a Resource Type Alert	46
6 Using Dynamic Groups and Policies	48
Working with Dynamic Groups	48
Working with Policies	51

- 7 Global Monitoring and Alerting Settings 54**
 - Configuring Global Alert Properties 54
 - Configure Alert Notification Throttling 55
 - Configuring Alert Notification Email Properties 55
 - Configure Hyperic Version and Security Announcements 56

- 8 Managing Alert and Notification Volume 57**
 - Preventing Multiple Alerts for the Same Problem 57
 - Disable all Alert Notifications 58
 - Configure Network Host Dependencies for Hierarchical Alerting 58
 - Setting a Notification Throttle 58

- 9 Manage Alert Definitions 59**
 - Viewing and Editing Alert Definitions 59
 - Enabling or Disabling Alert Definitions 61
 - Enable Alert Definition Change Logging 63

- 10 Creating and Managing Resource Groups and Applications 65**
 - Configuring and Managing Resource Groups 65
 - Defining Alerts for a Resource Group 68
 - Schedule Downtime for a Resource Group 70
 - Creating and Managing Applications 70

- 11 Map vCenter Virtual Machine IDs to vRealize Hyperic Platforms 73**

- 12 Configuring Hyperic as an SNMP Trap Receiver 74**
 - Step 1 - Select and Enable an Agent to Receive Traps 74
 - Step 2 - Configure Trap Generator 75
 - Step 3 - Create a Proxy Platform for Trap Data 75
 - Resolving Configuration Issues 76

- 13 Managing Hyperic Users and Roles 77**
 - Creating and Manage User Accounts 77
 - Creating and Managing Roles in vRealize Hyperic 80
 - Understanding Role-based Dashboards 86

- 14 View and Manage Resource Data 88**
 - Navigate to the Inventory Page for a Platform 89
 - Viewing Inventory Properties for a Platform 89
 - Navigate to the Inventory Page for a Server 93
 - Viewing Inventory Properties for a Server 93
 - Navigate to the Inventory Page for a Service 95

- Viewing Inventory Properties for a Service 95
- Navigate to the Inventory Page for a Group 97
- Viewing Inventory Properties for a Group 97
- Navigate to the Inventory Page for an Application 98
- Viewing Inventory Properties for an Application 99
- Creating Network Entities Manually 101
- Remove a Resource from the Inventory 108

15 Deploying and Managing Plug-ins 109

- Plugin Administration Tasks 110
- Resource Plugin Deployment Directories 110
- Working with Plugin Manager 110
- About Custom Plugin Deployment with Plugin Manager 111
- Removing a Plugin with Plugin Manager 111
- Synchronizing Plugins at Startup 112
- Disable Server-Agent Plugin Synchronization 113

About vRealize Hyperic Administration Guide

The *vRealize Hyperic Administration Guide* provides information about adding resources to VMware[®] vRealize™ Hyperic[®], and how to configure monitoring and alerts, and to manage users and roles.

Intended Audience

This information is intended for anyone who wants to install, upgrade, or use ESX. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

Discover and Import Resources to Inventory



There are various ways to discover and import resources to the vRealize Hyperic inventory.

- [Import or Skip Resources in Auto-Discovery Portlet](#)
This topic describes how to process the contents of the Auto-Discovery Portlet.
- [Using the Auto-Discovery Results Page](#)
This topic describes the Auto-Discovery Results page and its components.
- [Import or Skip Resources in Auto-Discovery Results Page](#)
This topic describes how to process the contents of the Auto-Discovery Results page.

Import or Skip Resources in Auto-Discovery Portlet

This topic describes how to process the contents of the Auto-Discovery Portlet.

You can process the contents of the Auto-Discovery Portlet in these ways:

Prerequisites

If the vRealize Hyperic agent discovered 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. Note however, that some level of configuration is required to start managing some resources types - see the Configuration Properties section on a resource's Inventory tab for configuration requirements.

Procedure

1

Option	Description
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 .

2 To import selected resources

- De-select the resources you do not want to add to inventory, and click **Add to Inventory**,
Or
- De-select the resources you do want to add to inventory and click **Skip Checked Resources**.

If you do not import a resource displayed in the Auto-Discovery portlet, note:

- ■ If you skip a new platform, you skip its servers as well.
 - During the next platform scan, skipped resources will reappear in the portlet after the next scan that detects them.

Using the Auto-Discovery Results Page

This topic describes the Auto-Discovery Results page and its components.

The Auto-Discovery Results page appears when you click a platform name in the Dashboard's Auto-Discovery portlet. The page shows the results of the most recent platform scan - the new or changed resource information that was discovered, as well as information for unchanged resources. You can filter the page for new, modified, or unchanged resources. You can selectively import new and changed data to inventory.

The Auto-Discovery Results page has three sections:

- Platform Type — This section contains the following data and controls

Platform Type	The resource type of the platform.
Fully Qualified Domain Name	The platform's FQDN
Import Values/Do Not Import	In this version of Hyperic, this control has no effect. Changing the value has no impact on what values are imported.
- Network Properties — This section contains the following data and controls:
 - All States/New/Modified/Unchanged selector — This pulldown allows you to filter the IP addresses that appear in the list by change state.
 - The following data is shown for each IP address matching the currently selected filter value:
 - IP Address
 - Netmask
 - MAC Address
 - Properties State - Value indicates if the connection properties are "New", "Unchanged", or "Changed".
- Servers — This section contains the following data and controls:
 - Server Types selector ---This pulldown allows you restrict the list to servers of a particular type.
 - All States/New/Modified/Unchanged filter ---This pulldown allows you to filter the servers that appear in the list by change state.
 - The following data is shown for each for each server in the list:
 - Server---The name of the resource.

- Server type---The resource type of the server.
- Install Path---Where the server is installed.
- Server Status---Value indicates whether the server is "New", "Unchanged", or "Modified".
- Import Values/Do Not Import (callout #5 in screenshot)---The value of this pull-down controls whether the server data will be imported when the OK button is clicked.

Import or Skip Resources in Auto-Discovery Results Page

This topic describes how to process the contents of the Auto-Discovery Results page.

If 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. For more information, see the Configuration Properties section on a resource's Inventory tab for configuration requirements.

To process the contents of the Auto-Discovery Results page:

Procedure

- 1 View the new and changed properties in the Network Properties section of the page.
- 2 If you do not want to import the new or changed network properties to the Hyperic inventory, select **Do Not Import** from the pulldown in the Platform Type section of the page
- 3 View the new and changed servers listed in the Servers section of the page.
- 4 If you do not want to import a new or changed server to Hyperic inventory, select **Do Not Import** from the pulldown in the Action column.
- 5 Click **OK** to import the resource data which you have approved for import.

Options for Running and Controlling Resource Discovery

2

This section describes options for initiating and preventing vRealize Hyperic resource discovery processes.

- [Configure the Frequency of the Default Scan](#)
This section describes how to configure the frequency of the default scan.
- [Configure the Frequency of the Runtime Scan](#)
This section describes how to configure the frequency of the runtime scan.
- [Scan a Platform On-Demand](#)
This section describes how to initiate a platform scan.
- [Solving Auto-Discovery Problems](#)

Configure the Frequency of the Default Scan

This section describes how to configure the frequency of the default scan.

The default scan detects servers and platform services, typically using the process table or the Windows registry. Default scans are less resource-intensive than runtime scans. By default, an agent runs a default scan upon startup and every 15 minutes thereafter. To change the frequency:

Procedure

- ◆ Uncomment the `autoinventory.defaultScan.interval.millis` line in the `agent.properties` file, and set the desired value in milliseconds.

Example: Disable Default Scan

To prevent the agent from running default scans on the platform, uncomment the `autoinventory.defaultScan.interval.millis` parameter in the `agent.properties` file, and set it to `-1`, as follows:

```
autoinventory.defaultScan.interval.millis=-1
```

Configure the Frequency of the Runtime Scan

This section describes how to configure the frequency of the runtime scan.

By default, an HQ Agent runs a runtime scan once a day. A runtime scan may use more resource-intensive methods to detect services than a default scan. For instance, a runtime scan may involve issuing an SQL query or looking up an MBean. To change the frequency of the runtime scan:

Prerequisites

Procedure

- ◆ Edit the `autoinventory.runtime.Scan.interval.millis` line in the `agent.properties` file, and set the desired value in milliseconds.

Scan a Platform On-Demand

This section describes how to initiate a platform scan.

To initiate a platform scan:

Prerequisites

You can only initiate an auto-discovery scan for a platform that runs an Hyperic Agent. Platforms that are remotely monitored, such as Cisco and other network device platforms that do not run an agent, cannot be scanned.

Procedure

- 1 Navigate to the platform you want to scan.
- 2 Select **Tools > New Auto-Discovery**.
- 3 You can run default scan only, or run a file scan in addition to the default scan.

Option	Description
Default Scan	Click OK at the top of the page.
File Scan	Follow the procedure below

- a Check the server types that vRealize Hyperic should look for on the platform.
- b In `ScanDirs`, specify the directories that vRealize Hyperic should scan.
- c In `ExcludeDirs`, specify the directories that vRealize Hyperic should not scan.
- d In `fsTypes`, select the type of file system to scan: local disks, network-mounted disks, or both (all).
- e In `depth`, specify the depth in the directory structure to which vRealize Hyperic should scan.

- f To have vRealize Hyperic follow symlinks when scanning, select **Should symlinks be followed**.
- g Click **OK**.

Solving Auto-Discovery Problems

Setting Up Alert Notifications and Escalations

3

- [Tailoring Alert Notification Templates](#)

The content and format of alert notifications are governed by Groovy server pages.

- [Configuring Roles for Role-Based Alert Notifications](#)

- [Enabling SNMP Trap Notifications](#)

This section has information about enabling vFabric Hyperic to send SNMP traps to an SNMP management system.

- [Enable Syslog Notifications](#)

This section describes how to enable the syslog notifications.

- [Configuring and Managing Escalations](#)

This section provides information about configuring and managing escalations.

- [Setting Up Alert Action Scripts](#)

This topic describes action scripts and how they can be used.

- [Configuring a Custom Control Action](#)

This section has instructions about how to configure vRealize Hyperic to run a script or executable that implements a resource control action.

- [Environment Variables for Fired Alert Data](#)

This topic lists the environment variables for fired alert data and their descriptions.

Tailoring Alert Notification Templates

The content and format of alert notifications are governed by Groovy server pages.

The pages are located in the `\server-n.n.n-EE\hq-engine\hq-server\webapps\ROOT\WEB-INF\alertTemplates` directory.

The following files are included.

Table 3-1. Alert Template Files

File Name	Purpose
subject.gsp	Defines the subject line in email notifications.
html_email.gsp	Defines the body of an HTML email alert notification.
text_email.gsp	Defines the body of a text email alert notification.
sms_email.gsp	Defines the body of an SMS alert notification.
opennms_notify.gsp	Defines an OpenNMS SNMP trap.
snmp_trap.gsp	Defines an SNMP trap.

Email Notification Recipients

You can send email notifications to individually name recipients, all users with a specific role, or to recipients who are not vRealize Hyperic users. Email to a vRealize Hyperic user is sent in the mail format specified in the user's user profile.

Email Notification Subject Line

The subject line of an email notification contains the string HQ followed by the alert priority, the alert definition name, the name of the resource for which the alert conditions were met, and the alert status. The format of the subject line is [HQ] `${priority}` – `${alertDef.name}` `${resource.name}` `${status}`. For example, [HQ]!! – High Disk Usage saint Win32 File System R:(remote/NTFS) notifies the recipients of a medium-high alert named "High Disk Usage" that was triggered for a resource named "saint Win32 File System R:\ (remote/NTFS)".

Email Notification Body

A variety of details are included in the body of an alert notification.

- Resource name
- Alert name
- Date and time of the alert
- Condition that triggered the alert. For example, the threshold value condition that triggers the alert and the specific value that triggered this alert.
- Severity of the alert
- In the case of an alert on an individual resource, the last value of the each of the resource's indicator metrics before the alert was triggered
- In the case of an alert for a group of resources, a list of the group's member resources that triggered the alert, and a link to a chart that displays the metrics that triggered the alert
- A link to the alert detail in the user interface

Available Variables for Notification Templates

There are a number of variables that you can reference in the `.gsp` templates that govern the content of alert notifications. You use alert-related classes to define the variables.

You insert variables in the format `$ClassVar.FieldVar`, where `ClassVar` is the class with the field, and `FieldVar` is a field for which the class contains a `get()` method. For example, the variable `{alertDef.description}` generates the alert definition description from an `AlertDefinitionInterface` class.

You append the variable for the required field to the variable of class that provides it, using a dot (`.`) separator to specify the variable for a field.

For data about...	from this class...	use this class variable	and a field variable...	resulting in...
The alert definition for the fired alert	<code>AlertDefinitionInterface</code>	<code>alertDef</code>	<code>description</code> <code>id</code> <code>name</code> <code>priority</code>	<code>alertDef.description</code> <code>alertDef.id</code> <code>alertDef.name</code> <code>alertDef.priority</code>
The fired alert	<code>AlertInterface</code>	<code>alert</code>	<code>id</code> <code>timestamp</code>	<code>alert.id</code> <code>alert.timestamp</code>
Alert actions	<code>ActionExecutionInfo</code>	<code>action</code>	<code>longReason</code> <code>shortReason</code>	<code>action.longReason</code> <code>action.shortReason</code>
Resource that fired the alert	<code>Resource</code>	<code>resource</code>	<code>instanceid</code> <code>mtime</code> <code>owner</code> <code>prototype</code> <code>resourceType</code>	<code>resource.instanceId</code> <code>resource.mtime</code> <code>resource.owner</code> <code>resource.prototype</code> <code>resource.resourceType</code>
Notification recipients	<code>AuthzSubject</code>	<code>user</code>	<code>active</code> <code>department</code> <code>emailAddress</code> <code>firstName</code> <code>fullName</code> <code>htmlEmail</code> <code>lastName</code> <code>phoneNumber</code> <code>smsAddress</code>	<code>user.active</code> <code>user.department</code> <code>user.emailAddress</code> <code>user.firstName</code> <code>user.fullName</code> <code>user.htmlEmail</code> <code>user.lastName</code> <code>user.phoneNumber</code> <code>user.smsAddress</code>
Status of an alert. Can be used only in <code>subject.gsp</code>		No class variable required	<code>status</code>	<code>status</code>
Whether an escalation action is an SMS notification. Can be used only in <code>subject.gsp</code>		No class variable required	<code>IsSms</code>	<code>IsSms</code>

Configuring Roles for Role-Based Alert Notifications

An authorized user can access the New Role command in the Administration tab of the HQ user interface. You can also create HQ roles using the HQApi role command.

Enabling SNMP Trap Notifications

This section has information about enabling vFabric Hyperic to send SNMP traps to an SNMP management system.

This section has information about enabling vFabric Hyperic to send SNMP traps to an SNMP management system.

Configure vRealize Hyperic Server to Send SNMP Traps

This section has information about enabling vRealize Hyperic to send SNMP traps to an SNMP management system.

Procedure

- 1 Click **HQ Server Settings** on the Administration page.
- 2 At the bottom of the page, in the SNMP Server Configuration Properties section, define the properties for your version of SNMP. For more information, refer to the SNMP tables respective to your version.

After the configuration above is complete, the SNMP Trap notification tab is available when you define or edit an alert definition.

vRealize Hyperic Server SNMP v1 Parameters

This section lists the parameters for SNMP v1, and their descriptions.

SNMP v1 Properties

Table 3-2. SNMP v1 Properties

Configuration Option	Description	Allowable Values
SNMP Trap OID	The OID of the notification to be sent. Supplies the value of <code>snmpTrapOID</code> . 0 - the second varbind in a trap or inform that Hyperic Server generates. (The first varbind is <code>SysUpTime.0</code> .)	
Default Notification Mechanism	Your selection governs the notification type that will appear as the default notification type option in the "Notification Mechanism" pulldown list that is presented in configuration dialogs when user configures an SNMP notification as an alert action, or as a step in an escalation.	For v1 of the SNMP protocol, choose V1 Trap. This is the only trap type you can generate for SNMP v1.

Table 3-2. SNMP v1 Properties (Continued)

Configuration Option	Description	Allowable Values
Enterprise OID	Enterprise OID	
Community	The community name to be sent with the trap.	
Generic ID	Single digit identifier of the trap type.	<ul style="list-style-type: none"> ■ 0 - coldStart ■ 1 - warmStart ■ 2 - linkDown ■ 3 - linkUp ■ 4 - authenticationFailure ■ 5 - egpNeighborLoss ■ 6 - enterpriseSpecific
Specific ID	The specific trap code for an enterprise-specific trap (when Generic ID is set to 6).	
Agent Address	Address of the managed object that generates the trap.	

vRealize Hyperic Server for SNMP v2c Parameters

This section lists the parameters for SNMP v2c, and their descriptions.

SNMP v2c Properties

Table 3-3. SNMP v2c Properties

Configuration Option	Description	Allowable Values
SNMP Trap OID	The OID of the notification to be sent. Supplies the value of <code>snmpTrapOID</code> . 0 - the second varbind in a trap or inform that Hyperic Server generates. (The first varbind is <code>SysUpTime.0</code> .)	
Default Notification Mechanism	Specifies the default notification type that will appear in configuration dialogs when an authorized user configures an SNMP notification as an alert action, or as a step in an escalation. This choice simply defines the default option - the user configuring an alert action or escalation can choose a different message type.	<ul style="list-style-type: none"> ■ V1 Trap ■ V2c Trap ■ Inform
Community	The community name to be sent with the trap.	

vRealize Hyperic Server for SNMP v3 Parameters

This section lists the properties for enabling vRealize Hyperic to send SNMP notifications to an NMS. When vRealize Hyperic is so enabled, you can use SNMP notifications in alert definitions - as alert actions and escalation steps.

SNMP v3 Parameters

Table 3-4. SNMP v3 Parameters

Configuration Option	Description	Allowable Values
SNMP Trap OID	The OID of the notification to be sent. Supplies the value of <code>snmpTrapOID</code> . 0 - the second varbind in a trap or inform that Hyperic Server generates. (The first varbind is <code>SysUpTime.0</code> .)	
Default Notification Mechanism	Specifies the default notification type that will appear in configuration dialogs when an authorized user configures an SNMP notification as an alert action, or as a step in an escalation. This choice simply defines the default option - the user configuring an alert action or escalation can choose a different message type.	<ul style="list-style-type: none"> ■ V1 Trap ■ V2c Trap ■ Inform
Security Name	The username Hyperic's SNMP agent should use when sending notifications to the NMS.	Required.
Local Engine ID	ID of Hyperic's SNMP agent; this value appears automatically, and is not user-configurable.	
Auth Protocol	The SNMP authentication protocol Hyperic Server should use for communications with the NMS.	<ul style="list-style-type: none"> ■ None ■ MD5 ■ SHA
Auth Passphrase	The SNMP authorization passphrase configured for use when communication with the NMS.	
Privacy Protocol	The SNMP Privacy Protocol Hyperic Server should use for communication with the NMS.	<ul style="list-style-type: none"> ■ None ■ DES ■ 3DES ■ AES-128 ■ AES-192 ■ AES-256
Privacy Passphrase	The SNMP privacy passphrase configured for use when communication with the NMS.	

Table 3-4. SNMP v3 Parameters (Continued)

Configuration Option	Description	Allowable Values
Context Engine ID	The EngineID of the NMS. This, along with Context Name, identifies the SNMP context for accessing management data.	Required for v1 and v2c traps. Do not supply for Inform.
Context Name	The name of the SNMP context that provides access to management information on the NMS. A context is identified by the Context Name and Context Engine ID.	

Enable Syslog Notifications

This section describes how to enable the syslog notifications.

You enable syslog notification in the HQ server's `server-log4j.xml` file, in the `SERVER_HOME/conf` directory.

Procedure

- 1 Open the `server-log4j.xml` file in a text editor.
- 2 Insert the `-r` switch to `syslogd` at startup to enable the syslog receiver.

This enables remote logging.

- 3 To enable the syslog appender, uncomment the following lines and change level value to `INFO` instead of `ERROR`

```
<appender name="SYSLOG" class="org.apache.log4j.net.SyslogAppender">
  <errorHandler class="org.apache.log4j.helpers.OnlyOnceErrorHandler"/>
  <param name="Facility" value="SYSLOG"/>
  <param name="FacilityPrinting" value="true"/>
  <param name="SyslogHost" value="localhost"/>
  <layout class="org.apache.log4j.PatternLayout">
    <param name="ConversionPattern" value="%c{1}[%r]: %m%n"/>
  </layout>
</appender>

<logger name="org.hyperic.hq.bizapp.server.action.log.SyslogAction">
  <level value="INFO" />
  <appender-ref ref="SYSLOG" />
</logger>
```

If you want to configure a different facility for the appender, note that it must be syslog-configured. If `server-log4j.xml` does not contain the lines shown above, add them in the appenders section. All appenders in the file must be grouped together.

What to do next

To present the syslog notification as an option when you set up an escalation, restart the vRealize Hyperic server.

Configuring and Managing Escalations

This section provides information about configuring and managing escalations.

Create an Escalation

This section describes the steps necessary to create an escalation.

Procedure

1 [Create an Escalation Scheme](#)

This section describes how to create an escalation scheme.

2 [Create Escalation Actions](#)

This section describes how to create an escalation action for the escalation scheme.

Procedure

- 1 Create an escalation scheme.
- 2 Create escalation actions.

Create an Escalation Scheme

This section describes how to create an escalation scheme.

Procedure

- 1 In the masthead, click **Administration**.
- 2 Click **Escalation Schemes Configuration**.
- 3 Enter values in the **Name** and, optionally, **Description** fields.
- 4 Configure acknowledgment options in the If the alert is acknowledged section:
 - Allow user to pause escalation for - Click to enable a user to pause the escalation when acknowledging the alert. Select **Until Fixed** or a duration from the pulldown list. (Options range from 5 minutes to 72 hours.) A user acknowledging an alert with this escalation will have the option of pause the escalation process for the period you specify.
 - Continue escalation without pausing - With this default value, a user acknowledging an alert with this escalation will not be offered the option to pause the escalation.

- 5 Configure state change notification options in the If the alert state has changed section:
 - Notify previously notified users of the change - With this default setting, when the state of the alert changes, state change notifications will be sent only to recipients who have already received a notification in previous escalation steps.
 - Notify entire escalation of the change - Click if you want alert state change notifications to be sent to every notification recipient in the escalation - whether or not they have received a previous notification.
- 6 Configure escalation repeat behavior in the "If alert is fixed" section:
 - Stop escalation execution - With this default setting, the escalation will not be repeated for an alert that is unfixed at the end of the escalation.
 - Repeat escalation actions - Click to repeat the escalation process if the alert has not been fixed by the end of the escalation.
- 7 Click **Next Step**.

What to do next

[Create Escalation Actions](#)

Create Escalation Actions

This section describes how to create an escalation action for the escalation scheme.

Prerequisites

Verify that you have completed the instructions in [Create an Escalation Scheme](#).

Procedure

- 1 On the Escalation Configuration page, click **Create Action**.
- 2 In the Create Escalation Scheme Actions section, select an action type and proceed to the directions for that type.
 - Email. For more information, see [Create an Email Action](#)
 - SMS. For more information, see [Create an SMS Action](#)
 - Sys Log. For more information, see [Create a Sys Log Action](#)
 - SNMP Trap. For more information, see [Create an SNMP Notification Action](#)
 - Suppress Alert. For more information, see [Create a Suppress Alerts Action](#)

Create an Email Action

This section provides instructions for creating an email action.

Procedure

- 1 From the **Choose Who to Notify** menu, select a notification target type:
 - Notify Roles - You are prompted to select one or more roles. The users with those role assignments will be notified.
 - Notify HQ Users - You are prompted to select one or more Hyperic users to be notified.
 - Notify Other Recipients - You are prompted to enter a comma-separated email list for the individuals you want to notify.
- 2 Select **Continue** if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pulldown. The intervals range from 5 minutes to 24 hours.
- 3 Click **Save**.
- 4 Repeat this procedure to add another step to the escalation.

Create an SMS Action

After creating an action type, you must specify the action type.

Procedure

- 1 From the **Choose Who to Notify** menu, select a notification target type:
 - Notify Roles - You are prompted to select one or more roles. The users with those role assignments will be notified.
 - Notify HQ Users - You are prompted to select one or more Hyperic users to be notified.
 - Notify Other Recipients - You are prompted to enter a comma-separated email list for the individuals you want to notify.
- 2 Select **Continue** if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pulldown. The intervals range from 5 minutes to 24 hours.
- 3 Click **Save**.
- 4 Repeat this procedure to add another step to the escalation.

Create a Sys Log Action

After enabling notifications for escalations, you can create a sys log action.

Prerequisites

To enable this option, verify that you have completed the steps in [Enable Syslog Notifications](#).

Procedure

- 1 Supply the value for these segments of the syslog message:
 - meta
 - product
 - version
- 2 Select **Then continue** selected if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pulldown. The intervals range from 5 minutes to 24 hours.
- 3 Click **Save**.
- 4 Repeat this process to add another step to the escalation.

Create an SNMP Notification Action

This section describes how to create an SNMP notification action.

You can define an SNMP notification to be performed as a step in an escalation if the vRealize Hyperic server is configured for your NMS.

The trap or inform sent when the escalation step is performed contains three variable bindings:

- sysUptimeOID.0 - No configuration is required for this binding.
- snmpTrapOID.0 - This binding is configured on the HQ Server settings page. You can customize this variable for a specific action.
- A variable binding for the alert data specified in `snmp_trap.gsp`, a Groovy Server Page template that returns the alert definition name and the "short reason" for firing. This template can be customized, as desired. [Tailoring Alert Notification Templates](#).
- any additional variable bindings you define.

To configure an SNMP notification as an escalation step:

Procedure

- 1 After selecting the "SNMP Notification" action type, enter the following details:
 - a In the **IP Address** field, enter the address and port of the target SNMP server.
 - b In the **Use Default SNMP Trap OID** field,
 - Select the checkbox to use the default.
 - Clear the checkbox to specify a different SNMP Trap OID.
 - c In the **Notification Mechanism** field, select a notification type to send.
 - d In the **OID** field, enter the OID of the notification that will contain the alert details specified in the `snmp_trap.gsp`, template.

- 2 To add additional variable bindings, click **Add Another Variable Binding** and enter
 - **OID** - Enter an additional OID to include in the notification.
 - **Value** - Enter a value for the OID. You can enter plain text, or an alert variable.
- 3 Select **Then continue** if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pull-down. The intervals range from 5 minutes to 24 hours.
- 4 Click **Save**.
- 5 Repeat this procedure to add another step to the escalation.

Create a Suppress Alerts Action

This section describes how to create a suppress alerts action.

This action stops the alert from repeated firing - it is useful if the alert definition for the fired alert is not configured to "fire once until fixed", and but you want to stop repetitive firing at a specific point in the escalation process. After this step in the escalation process is performed, the alert will not fire again until fixed.

After selecting the "SNMP Notification" action type:

Procedure

- 1 Select **Then continue** if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pull-down. The intervals range from 5 minutes to 24 hours.
- 2 Select **Then continue** if you want the next step in the escalation to occur immediately after the current one. To specify a delay before the next step is performed, select an interval from the pull-down. The intervals range from 5 minutes to 24 hours.
- 3 Repeat this procedure to add another step to the escalation.

Setting Up Alert Action Scripts

This topic describes action scripts and how they can be used.

A script action allows you to access and use vRealize Hyperic environment variables that contain information about a fired alert. Script actions execute one at a time. Until a script action completes, additional alerts will not fire on the resource. You can use the data in any fashion you wish, for instance in a web service call to external management system.

To use the script action feature, write a script that implements the action or logic you wish to perform with the alert-related environment variables. When you configure the alert, specify the script to be executed when the alert fires. The script is server-side only, meaning it must be accessible and executable by the same user running the HQ Server process. For information about agent-side scripts,

Script actions can be defined for resource alerts and resource type alerts. In this version of Hyperic, escalation schemes do not support script actions.

Script actions are forked by the the vRealize Hyperic Server's Java process. Most operating systems duplicate the Java process's memory before executing the new process; exact behavior varies by operating system. If the operating system does not provide for over-committing memory, script execution requires an amount of free memory equal to the amount of memory that the Hyperic Server's Java process consumes - otherwise, the script action will not run.

Configuring a Custom Control Action

This section has instructions about how to configure vRealize Hyperic to run a script or executable that implements a resource control action.

Configuring a Custom Control Action

This section has instructions about how to configure vRealize Hyperic to run a script or executable that implements a resource control action.

1 [Install the Control Action Executable](#)

This topic provides instructions for installing the control action executable.

2 [Configure Control Action as a Platform Service](#)

You must configure the means by which an authorized user can invoke the custom control action.

Install the Control Action Executable

This topic provides instructions for installing the control action executable.

Install the executable that implements the control action on each managed platform where it will be used.

Prerequisites

The account under which the vRealize Hyperic agent runs must have:

- Adequate permissions to run the file.
- Adequate permissions to perform the tasks defined in the executable.

You can control the permissions required to run the file with `chmod` or an equivalent utility in your environment. If the tasks defined in the executable require permissions not granted to the account the Hyperic Agent runs under, one option is to add specific "NOPASSWD" entries in the `sudoers` file on the platform.

Procedure

- ◆ Create a "control_scripts" (or similar) directory in `/opt/hyperic`, or a directory high enough in the agent installation that it will not be overwritten during an agent upgrade.

Configure Control Action as a Platform Service

You must configure the means by which an authorized user can invoke the custom control action.

To do so, you configure it as a platform service of FileServer File Service type on each platform where the action will be used.

Procedure

1 Browse to the platform in the Resource Hub.

2 Select **Tools > New Platform Service**.

3 Name the service.

Including the platform name in the name for the new service will help you distinguish it from other services of the same type.

4 (Optional) In the **Description** field, describe the functionality of the control action.

5 Select **FileServer File** from the Service Type pulldown.

6 Click **OK**.

The Inventory page for the new platform service appears.

7 Click **Edit** in the Configuration Properties section of the Inventory page.

The Configuration Properties page for the new platform service appears.

8 On the Configuration Properties page, provide the following information:

Table 3-5. Configuration Properties

Parameter	Description
path	Enter the path to the executable in the path field in the Shared section, including the name of the file itself. This can be a relative path if the executable is in the vRealize Hyperic agent directory structure. A better practice is to store the script external to the agent directory structure, and specify the full path to the executable, for example, <code>/opt/hyperic/control_scripts/test.sh</code> .
timeout	Check the value in the timeout field in the Control section. This is the time in seconds HQ will wait for a response from the control before declaring it as a failed action. The default is 30 seconds, and under most circumstances should be adequate.

9 If the executable requires higher privileges than those available to the account running the Hyperic Agent, you can use the "prefix" field in the "Control" section to specify the sudo command. In this case, the local sudoers file has an entry for the user to run the script with a NOPASSWD directive.

10 Click **Ok**.

Environment Variables for Fired Alert Data

This topic lists the environment variables for fired alert data and their descriptions.

Environment Variables for Fired Alert Data

The environment variables for fired alert data are prefixed with the string "HYPERIC_". The table below describes the variables.

Table 3-6. Environment Variables for Fired Alert Data

Variable	Description	Example Output
HYPERIC_ALERT_ALERT_TIME	The time at which the alert fired, in milliseconds from epoch.	1219167000000
HYPERIC_ALERT_CONDITION	The condition that caused the alert to fire. Note This environment variable is supported only on Unix-based platforms. The value contains the Java \n character, which causes errors under Windows. In Windows environments, use the HYPERIC_ALERT_SUMMARY variable, which provides the same information without the \n character.	If Load Average 5 Minutes > 1.0 (actual value = 1.4)
HYPERIC_ALERT_DESCRIPTION	The description of the alert that fired.	This alert will fire when the load rises
HYPERIC_ALERT_ID	The internal Hyperic ID for the alert that fired. Note The HYPERIC_ALERT_ID for an alert is not committed to the Hyperic database until all alert actions are complete. Therefore, an alert action script (whether it uses SQL or HCApi) cannot query or update the Hyperic database using the alert's HYPERIC_ALERT_ID, because that value will not yet exist in the Hyperic database.	
HYPERIC_ALERT_NAME	The name of the alert that fired.	High Load
HYPERIC_ALERT_PRIORITY	The priority of the alert that fired, 1 for High, 2 for Medium, 3 for Low.	2
HYPERIC_ALERT_SUMMARY	A condensed data string that contains the relevant alert and resource names and values which triggered the alert.	Mac OS X DOWN The-Idea-Men Availability (0.0%)
HYPERIC_FIXED_ALERT_ID	Valid for recovery alerts only. Supplies the internal Hyperic ID for the primary alert to which the recovery alert is assigned.	

Table 3-6. Environment Variables for Fired Alert Data (Continued)

Variable	Description	Example Output
HYPERIC_FIXED_ALERT_NAME	Valid for recovery alerts only. Supplies the name of the primary alert to which the recovery alert is assigned.	High Load
HYPERIC_PLATFORM_NAME	The platform on which this alert fired.	localhost.hyperic.com
HYPERIC_RECOVERY_ALERT	A boolean that indicates if the alert is a recovery alert.	False
HYPERIC_RESOURCE_ID	The internal ID for the resource for which the alert fired.	
HYPERIC_RESOURCE_NAME	The name of the resource for which the alert fired.	localhost.hyperic.com

The following script is an example of using the alert variables. The example script simply writes the time that the script executed and the variables to a log file; it is not a representative use case. The purpose of script actions is to enable more complex alert actions, and actions that meet the unique needs of your environment.

```
#!/usr/bin/perl
my $logfile = "/tmp/output.txt";
my $date = localtime();
open LOGFILE, ">$logfile" or die "Cannot open log file for writing";
print LOGFILE "# Running script at $date", "\n";
foreach $key (sort keys(%ENV)) {
  if ($key =~ m/^HYPERIC/) {
    my $msg = "$key = $ENV{$key}";
    print LOGFILE $msg, "\n";
  }
}
close LOGFILE;
```

Configure Monitoring and Alerting for a Resource

4

There are several aspects to configuring monitoring and alerting for a resource.

- [Configure Metric Baselines](#)
This topic describes how to configure metric baselines.
- [Setting Up Log Tracking for a Resource](#)
IT problems can often be detected or diagnosed from messages generated by operating systems, application servers, network services, or middleware throughout the environment.
- [Setting Up Configuration Tracking for a Resource](#)
This section provides information about setting up configuration tracking for a resource.

Configure Metric Baselines

This topic describes how to configure metric baselines.

You configure the properties that control baseline calculations on the Server Settings page, available on the Administration tab.

Procedure

- ◆ In vRealize Hyperic, configure the following properties to control the baselining process. Changing the data set used to calculate baselines can affect baseline accuracy.

Table 4-1. Global Baselining Properties

Server Setting	Description	Default
Baseline Frequency	The frequency with which vRealize Hyperic calculates a baseline for each metric.	3 Days
Baseline Dataset	The time range of metric data used in calculating the baseline.	7 days

Table 4-1. Global Baseline Properties (Continued)

Server Setting	Description	Default
Baseline Minimum Data Points	The minimum number of data points used in calculating a baseline.	40
Track Out-of-Bounds Metrics	Controls whether or not Hyperic tracks out-of-bounds metrics - measurements that are greater than expected high range for a metric, or less than the expected low range for a metric.	off

Recalculate Baselines for a Resource

This topic describes how to recalculate baselines for a resource.

The baseline and expected high and low range values for every metric are calculated periodically — by default, every three days.

For an individual resource, you can recalculate baseline and high/low range values on-demand; when you do, the calculation will be based on the currently selected metric display range. The next time vRealize Hyperic does its periodic baseline calculations, baselines will again be based on the configured "Baseline Dataset" property.

To recalculate the metric baseline and expected range values for one or more metrics for a resource.

Procedure

- 1 Navigate to the **Metric Data** tab on the resource's Monitor page.
- 2 Set the metric display range upon which you wish to calculate baselines.
- 3 Select each metric whose baseline values you want to set.
- 4 Click **Set Baselines**.

The Low Range, Baseline, and High Range values for the selected metrics will be set to the low, average, and peak values for the current display range.

You can also initiate re-calculation of the baseline values for a single resource metric when you chart the metric.

Reset Baselines for a Charted Metric

This topic provides instructions about how to recalculate baselines for an individual metric from a metric chart.

Procedure

- 1 Chart the metric whose baseline you want to recalculate — click the chart icon next to it on the **Metric Data** minitab.

- 2 Click **Edit Range**, select the desired period of history upon which to base the baseline, and redraw the chart.
- 3 In the Metric Baseline & Expected Range section, click **Change Value** next to the **Baseline** field. A new baseline value is displayed. For example:

```
New value calculated from current Display Range: 26.5 MB
```

- 4 Click **Save Value** to save the calculated value as the baseline value for the metric. New values for High Range and Low Range appear.
- 5 To modify the values for High Range and Low Range:
 - a Click **Change Value** next to the **High Range** and **Low Range** field.
 - b Enter a value.
 - c Click **Save Value**.

Setting Up Log Tracking for a Resource

IT problems can often be detected or diagnosed from messages generated by operating systems, application servers, network services, or middleware throughout the environment.

vRealize Hyperic can monitor messages in log files and in memory, and record events in the vRealize Hyperic database based on criteria you specify.

- [Configuring Log Tracking](#)
This topic provides information about the supported resource types and message types, as well as configuration options.
- [View Log Events](#)
Log events for a particular resource are indicated in the timeline at the bottom of the resource's Indicators page.
- [Log Tracking Support Classes](#)
This topic provides information about the different support classes.

Configuring Log Tracking

This topic provides information about the supported resource types and message types, as well as configuration options.

Hyperic Resource Types that Support Log Tracking

Hyperic supports log tracking for operating system platforms, network services, and most server types. If a resource supports log tracking, its Configuration Properties page contains log tracking configuration options.

Supported Log Message Types

Hyperic can monitor and record events for:

- Log file messages that specify log levels using log4j log levels.
- Events written to Windows Event Logs.
- Network request results for a variety of network services.

Log Tracking Configuration Options

You enable and configure log tracking for a resource on its Configuration Properties page. Navigate to the resource's Inventory page, and click Edit in the Configuration Properties section to display the Configuration Properties page.

Note Log and configuration tracking must be enabled for a resource if you wish to log events for log messages or configuration changes. Event logging is automatic for alerts and control actions. Log tracking configuration options vary somewhat by resource type.

Log Tracking for Resources with Log4j Logs

This section describes the log tracking configuration options for resources that whose log files use logj4 levels.

An authorized user can set the values of these configuration options:

- Enable/disable log tracking.
- Specify one or more files to track, as a comma-separated list. The Hyperic Agent must be able to read these files, so make sure permissions are set appropriately.
- Specify the highest log level to track:
 - Error - Messages with log level FATAL or ERROR
 - Warn
 - Info
 - Debug
- Specify substrings or expressions to use as include/exclude filter criteria. Enter a substring or a regular expression that a log message must contain or match in Log Pattern Match.

Log Tracking for Network Services

This section describes the log tracking configuration options for network services.

An authorized user can set the values of these configuration options:

- Enable/disable log tracking.

- Specify the highest log level to track:
 - Error
 - Warn
 - Info
 - Debug
- Specify substrings or expressions to use as include/exclude filter criteria. Enter a substring or a regular expression that a log message must contain or match in Log Pattern Match.

Log Tracking for Windows Platforms

This section describes the log tracking configuration options for platforms of type "win32".

An authorized user can set the values of these configuration options:

- Enable/disable log tracking.
- Specify one or more Event Log to track:
 - System - contains events logged by Windows system components. For example, if a driver fails to load during startup, an event is recorded in the system log. Windows predetermines the events that are logged by system components.
 - Application - contains events logged by programs. For example, a database program may record a file error in the application log. Events that are written to the application log are determined by the developers of the software program.
 - Security - contains events such as valid and invalid logon attempts, as well as events related to resource use, such as the creating, opening, or deleting of files. For example, when logon auditing is enabled, an event is recorded in the security log each time a user attempts to log on to the computer. A Windows administrator or member of the Windows Administrators group specify which events are recorded in the security log.
 - * - causes all event logs to be tracked.
- Specify the highest log level to track:
 - Error
 - Warn
 - Info
 - Debug

Content of Logged Windows Events

When Windows log tracking is enabled, an entry of this form is logged for events that match the criteria you specified on the resource's Configuration Properties page:

```
[Timestamp] Log Message (EventLogName):
EventLogName:EventAttributes
```

where:

- Timestamp - is when the event occurred
- Log Message - is an text string
- EventLogName - is the Windows event log type, "System", "Security", or "Application".
- EventAttributes - a colon delimited string made of the Windows event Source and Message attributes.

For example, this log entry:

```
04/19/2010 06:06 AM Log Message (SYSTEM): SYSTEM: Print:
Printer HP LaserJet 6P was paused
```

is for an Windows event written to the Windows System event log at 6:06 AM on 04/19/2010. The Windows event Source and Message attributes, are "Print" and "Printer HP LaserJet 6P was paused.", respectively.

Tailoring the Content and Format of Logged Windows Events

You can configure the last portion of the log data that the agent writes for a Windows event - referred to above as EventAttributes. You can include additional event attributes, for example User and Computer. To do so, you add the `platform.log_track.eventfmt` property to the `agent.properties` file for the vRealize Hyperic agent monitoring the Windows platform.

platform.log_track.eventfmt Property

This property specifies the content and format of the Windows event attributes that a vRealize Hyperic agent includes when logging a Windows event as an event in vRealize Hyperic.

`agent.properties` does not contain the `platform.log_track.eventfmt` property, you must explicitly add it if you want to tailor the data logged for Windows events.

When Windows log tracking is enabled, an entry of this form is logged for events that match the criteria you specified on the resource's Configuration Properties page.

You can use the parameters below to configure the Windows event attributes that the agent writes for a Windows event. Each parameter maps to Windows event attribute of the same name.

- `%user%` - The name of the user on whose behalf the event occurred.
- `%computer%` - The name of the computer on which the event occurred.
- `%source%` - The software that logged the Windows event.
- `%event%` - A number identifying the particular event type.
- `%message%` - The event message.
- `%category%` - An application-specific value used for grouping events.

View Log Events

Log events for a particular resource are indicated in the timeline at the bottom of the resource's Indicators page.

A circular indicator over the timeline indicates a timeslice in which one or more events of any type - log events, configuration change events, or alerts - were logged.

Procedure

- ◆ Click the event indicator to view the data collected at that time.

You can use the Event Center to view events over time for all, or selected groups of resources, and filter by log event severity.

Log Tracking Support Classes

This topic provides information about the different support classes.

For information about the vRealize Hyperic support class for log events for log file messages that specify a log4j level, see the section about the Log4JLogTrackPlugin.

For information about the vRealize Hyperic support class for tracking Windows Event logs, see Win32EventTrackPlugin.

Plugin classes monitor network services log events using the `plugin.getManager().reportEvent` method.

Setting Up Configuration Tracking for a Resource

This section provides information about setting up configuration tracking for a resource.

- [Understanding vRealize Hyperic Configuration Tracking Functionality](#)

For supported resource types, you can configure vRealize Hyperic to track changes made to configuration files (or other local files of interest) and log events for the associated resource. Configuration tracking is supported for most platform and server types; typically not for services.

- [Set Up Event-Based Configuration Tracking for a Server](#)

The instructions in the section apply to server types whose plugin uses event-based configuration tracking, listed below.

- [Set Up Polling-Based Configuration Tracking](#)

For platforms and services that support configuration tracking, and server types that have the agent uses a cryptographic hash function to continuously compare a original version of the file with the current version to see if it has changed.

Understanding vRealize Hyperic Configuration Tracking Functionality

For supported resource types, you can configure vRealize Hyperic to track changes made to configuration files (or other local files of interest) and log events for the associated resource. Configuration tracking is supported for most platform and server types; typically not for services.

Configuration tracking functionality for server types varies by resource type, plugins that have been updated to use the event-based tracking added in vRealize Hyperic 4.6.5 track more information about file changes than those that still use the still-supported polling-based tracking.

The plugins that are updated to use event-based tracking are listed in [Set Up Event-Based Configuration Tracking for a Server](#). Follow the instructions in that section for those server types. For other server types, and for platform and service types that support configuration tracking see [Set Up Polling-Based Configuration Tracking](#).

Set Up Event-Based Configuration Tracking for a Server

The instructions in the section apply to server types whose plugin uses event-based configuration tracking, listed below.

To enable and configure configuration tracking for a resource of one of the following server types:

- Tomcat
- Apache
- WebSphere
- WebLogic Server
- JBoss
- PostgreSQL
- MySQL
- Oracle

Procedure

- 1 Navigate to the server instance.
- 2 Select **Configure Server** from the **Tools** menu.
- 3 In the Monitoring section of the page:
 - a Verify that the **Enable Config Tracking** checkbox, which enables and disables configuration tracking for the server, is checked.
 - b In the Configuration Files field, specify the files you wish to monitor, as follows:
path; recurse; regex1| regex2, path2; recurse; regex1| regex2| . . . | regexN
 - c Click **OK**.

Example: Configuration Files Setting

For MySQL, the default value of the Configuration Files field, which specifies which files to track is:

```
/etc;false;my\.cnf;,/etc/mysql;false;my\.cnf;,,;false;my\.cnf;
,~/;false;\.my\.cnf;,/usr/local/mysql/data;false;my\.cnf;
,/usr/local/var;false;my\.cnf;,%WINDIR%;false;my\.ini|my\.cnf;
,c:;false;my\.ini|my\.cnf;,,;false;my\.ini|my\.cnf;
```

When configuration tracking is enabled, this filter causes vRealize Hyperic to track `my.cnf` in the following directories:

- MySQLHome/etc
- MySQLHome/etc/mysql
- MySQLHome
- the home directory of the user under which MySQL runs
- `/usr/local/mysql/data`
- `/usr/local/var`
- the Windows directory
- `c:/`

and track `my.ini` in the following directories:

- the Windows directory
- `c:/`
- MySQLHome

Set Up Polling-Based Configuration Tracking

For platforms and services that support configuration tracking, and server types that have the agent uses a cryptographic hash function to continuously compare a original version of the file with the current version to see if it has changed.

To enable and configure configuration tracking for a platform or service:

Procedure

- 1 Navigate to the resource instance.
- 2 Select **Configure ...** from the **Tools** menu.

3 On the Edit Configuration Properties page:

- a Select the **Enable Config Tracking** checkbox to toggle configuration tracking on and off for the resource.
- b Enter the path, relative to the resource's installation directory, to each file you wish to track. Separate files with a comma.
- c Click **OK**.

Configuring Monitoring and Alerting for a Resource Type

5

This section provides information about configuring monitoring and alerting for a resource type.

- [Tailoring Metric Collection for a Resource Type](#)
Metric collection settings for a resource type are configured on the Monitoring Defaults page for the resource type.
- [Define a Resource Type Alert](#)
This topic explains how to define an resource type alert.
- [Define a Recovery Alert for a Resource Type Alert](#)
A recovery alert is special type of alert definition that you pair with a properly configured primary alert definition to streamline alert management.

Tailoring Metric Collection for a Resource Type

Metric collection settings for a resource type are configured on the Monitoring Defaults page for the resource type.

An authorized user can tailor metric collection settings on a per resource basis (on the resource's Metric Data minitab in the Resource Hub.) When you save changes to the metric collection settings for a resource type on the Monitoring Defaults page, the settings you configure will apply to all resources of that type in inventory. So, any custom metric collection configuration for resources of that type will be overwritten.

However, if a user has chosen a different set of indicator metrics for a resource instance (on the resource's Indicators minitab in the Resource Hub) changing the indicator metrics for a resource type on the Monitoring Defaults page will not override the user's selections.

- [Navigate to the Monitoring Defaults Page for a Resource Type](#)
This topic provides instructions for accessing the Monitoring Defaults page for a resource type.
- [Manage Metric Collection](#)
This topic provides instructions for enabling and disabling metric collection, as well as changing the collection interval.
- [Set Indicator Metrics](#)
This topic provides instructions for setting indicators for metrics.

- [APIs for Metric Collection Settings](#)

This topic provides information about using APIs for metric collection.

Navigate to the Monitoring Defaults Page for a Resource Type

This topic provides instructions for accessing the Monitoring Defaults page for a resource type.

Procedure

- 1 Click the **Administration** tab.
- 2 Click **Monitoring Defaults** in the Server Settings section of the **Administration** tab.
- 3 Scroll to the desired resource type on the Monitoring Defaults page, and click **Edit Metric Template** in that row.

Manage Metric Collection

This topic provides instructions for enabling and disabling metric collection, as well as changing the collection interval.

Procedure

- 1 Collection of the metric is now enabled by default with the specified collection interval on every resource of this type. To disable the collection of a metric, select the checkbox next to the metric name at the left and click **Disable Collection** at the bottom of the page.
- 2 To change a metric collection interval:
 - a Select the checkbox to left of the metric name.
 - b In Collection Interval for Selected at the bottom of the page, enter a time value and select a unit of time measure for the collection interval.
 - c Click the **Collection Interval for Selected** control at the bottom of the Monitoring Defaults page.

Set Indicator Metrics

This topic provides instructions for setting indicators for metrics.

To select the indicator metrics for a resource type (indicator metrics are the metrics that are charted on a resource's Indicators minitab in the Resource Hub):

Procedure

- ◆ Select the checkbox to the left of each metric name you wish to be an indicator, and click **Set Selected Metrics as Indicators** at the bottom of the Monitoring Defaults page.

The metrics you configure will be the default indicators for the resource type.

Changes to indicator metrics apply to existing resources but only for users that have not explicitly changed the default indicator page on the resources (of the specified type). If a user has changed the default indicator page for a resource, that will not be overwritten with changes made here.

APIs for Metric Collection Settings

This topic provides information about using APIs for metric collection.

As an alternative to the updating updating metric collection settings from the user interface, you can use vRealize Hyperic APIs to perform updates from the command line or from scripts to perform bulk updates:

- `HQApi metric` command - for listing and updating metric collection settings for an individual resource.
- `Api metricTemplate` command - for listing and updating metric collection settings for an all instances of a resource type.

Define a Resource Type Alert

This topic explains how to define an resource type alert.

A resource type alert is applied to all resources of the specific type that currently exist and new resources of the type that get created in the future. Only administrators can create resource type alerts.

- [Understanding Alert Condition Sets](#)

This topic provides information about the structure of a condition set.

- [Understanding Alert Actions](#)

This topic provides information about the different aspects of an alert action.

Procedure

- 1 Select the target resource type.
 - a Click **Administration** in the masthead.
 - b Click **Monitoring Defaults** in Server Settings section of the page.
 - c On the Monitoring Defaults Configuration page, click **Edit Alerts** for the resource type for which you want to define an alert.

The Monitoring Defaults page will display any alert definitions already assigned to the alert.
 - d Click **New**.
- 2 Define the alert properties.
 - **Name** — Name assigned by the user creating an alert definition. A fired alert is identified, in the user interface and alert notifications, by the alert definition name and a timestamp. An alert definition name should clearly communicate the nature of the problem. For example, Down for an alert on availability, or Low Memory for an alert on free memory.
 - **Description** — Description entered by the user creating the alert definition.
 - **Priority** — The severity of the problem, as defined by the person creating the alert definition: Low, Medium, or High. A consistent policy for defining an alert definition priority makes it easier to triage problems appropriately. An alert's priority is shown in vRealize Hyperic pages that present alert status and in alert notifications. You can sort alerts by priority in vRealize Hyperic's Alert Center or Operations Center.

- Active — The current enabled/disabled status of the alert definition. Alerts only fire for enabled alert definitions. When an alert definition is disabled, Hyperic does not evaluate its condition or fire alerts for it.
- 3 Define the Alert Condition set and click **OK**.

For more information, refer to [Understanding Alert Condition Sets](#).
- 4 Define alert actions. For more information, refer to [Understanding Alert Actions](#).

Understanding Alert Condition Sets

This topic provides information about the structure of a condition set.

Condition Set

An alert condition specifies a resource metric value or event that will initiate the alert firing process.

The condition types you can choose when you define a alert vary by resource type and Hyperic version. If a condition type is not supported by your version of Hyperic or is not valid for the target resource, it will not appear as an option.

To define a condition, choose one of the following condition types, and supply required parameter values.

- Metric condition - To base the alert on the value of a metric that vRealize Hyperic collects for the resource:
 - Metric - Select a metric from the selector list. Only currently enabled metrics are listed.
 - Define the rule for evaluating the metric value. You can:
 - Compare metric value to an absolute value. Select an operator: >(greater than), <(less than), =(equal to), or != (not equal to), and enter a metric value. If the metric value is a percentage, specify it as a float value. For example, enter .99 for 99%, 1.0 for 100%. Use a period (.) as a decimal separator, rather than a comma (,).
 - Compare metric value to its minimum, baseline, or maximum value, in vRealize Hyperic only. Select an operator: >(greater than), <(less than), =(equal to), or != (not equal to), and choose Min Value, Baseline Value or Max Value. Baselining must be enabled.
 - Fire upon change in metric value. Click **value changes**.
- Inventory Property Condition - To define a condition that is triggered when the value of an inventory property for resource changes, select an inventory property. The pulldown contains only those inventory properties that are valid for the type of the resource to which the alert applies.
- Control Action Condition - When you define an alert for a resource that supports control actions, you can define a condition that is triggered when a particular control action is performed. If desired, you can base the condition on a control action with a particular result status: in progress, completed, or failed. Pulldowns allow you to select a control action that the resource supports, and a result status if desired.

- **Events/Log Level Condition** - To define a condition that is triggered by a log event, select a message severity level (error, warn, info, debug, all) and optionally a match string. The condition is satisfied each time a message of the selected severity that contains the match string (if one was specified) is written to a log file that vRealize Hyperic is tracking. Log tracking must be enabled for the resource. To determine the log files that vRealize Hyperic monitors for the resource, see the Configuration Properties section of the resource's Inventory tab. The log files that vRealize Hyperic monitors for a resource are defined using the `server.log_track.files` property.
- **Config Changed... Condition** - This type of condition is triggered by a change to a configuration file that vRealize Hyperic is configured to monitor for the resource. To limit the condition to a single file, enter its filename in the match filename field. If you don't specify a filename, a change to any file monitored will trigger the alert. To determine the log files that vRealize Hyperic monitors for the resource, see the Configuration Properties section of the resource's Inventory tab. The files that vRealize Hyperic monitors for a resource are defined using the `server.config_track.files` property. The maximum length for filename entered is 25 characters.

Define Additional Conditions

In vRealize Hyperic, you can define up to three conditions for an alert. To add another condition, click **Add Another Condition** and specify whether both the new condition and the preceding one must be satisfied for the alert to be triggered ("AND") or only one must be satisfied ("OR").

Define Recovery Alert Behavior

To designate the alert you're defining as a recovery alert, select the primary alert definition from the pulldown.

A recovery alert condition should detect when the condition that fired the primary alert is no longer true. When a recovery alert fires, it marks the primary alert Fixed, and the primary alert definition is re-enabled. The primary alert definition should be configured to generate one alert and then disable alert definition until fixed, as described below.

Enable Actions

You can make the condition absolute, or fire after the condition occurs repeatedly. Choose one of the following options:

- **Each time conditions are met.**

The alert fires upon a single occurrence of the condition.

- **Once every _ times conditions are met within a time period of _ minutes.**

This option configures an alert to fire when the condition(s) occur multiple times over a period of time. Enter the number of occurrences and period of time.

This condition depends on the time interval that you choose when setting the metric conditions.

The number of recurrences, multiplied by the time interval, must be equal or less than the "time period of minutes" that you specify. For example, if the available metric interval equals two minutes, and the condition is 3 recurrences in 6 minutes (2×3 is less than or equal to 6), the condition is valid. If the available metric interval equals two minutes, and the condition is 3 recurrences in 5 minutes, the condition is invalid and is not applied.

Enable Action Filters

An action filter can be used to control alert firing and alert actions.

Disable an Alert Definition Upon Firing

Click **Generate one alert** and then disable alert definition until fixed to disable the alert definition after firing and reenable it when the alert that triggered it is marked Fixed.

This option eliminates redundant firing for the same problem. If you do not choose this option, the alert will fire repeatedly as long as the triggering condition is still true.

This configuration option, used in conjunction with recovery alerts, automates the process of disabling and re-enabling an alert definition. The outcomes are that there are no redundant alerts for the same problem, and you do not have manually "fix" an alert triggered by a transient problem.

Disregard Control Actions for Related Alerts

The **Disregard control actions that are defined for related alerts** option appears on New Alert Definition pages for resources that support control actions. This option only applies when:

- The current alert definition will include an alert action.
- The resource associated with the alert is a member of an application.
- There are other members of the same application with alerts that fire control actions (ideally the same control action)

Under these circumstances, this configuration option ensures that if multiple alerts are fired within a short period for resources that are members of the same application, only one control action will be executed. For example, this would prevent a server from being restarted several times in a short period of time for the same alert conditions. For instance, you might have an alert with an action to restart a Tomcat server if the JVM Free Memory got too low and another alert with an action to restart the same server if the JVM Active Thread count got too high. If both alerts fired at the same time and they were filtering control actions, only 1 restart control action would be executed and not two.

Understanding Alert Actions

This topic provides information about the different aspects of an alert action.

You assign actions to an alert definition on the Alert Definition page, which appears when you save a new alert definition or edit an existing alert definition.

The Alert Definition page is similar to the New Alert page, with the addition of Edit controls in the Alert Properties and Condition Set sections, and tabs at the bottom of the page for defining alert actions.

You can specify multiple actions to be performed automatically when an alert fires. The types of actions available in the Alert Definition page vary based on: (1) the type of resource the alert applies to, (2) your version of vRealize Hyperic, and (3) whether you've configured vRealize Hyperic for the types of actions that must be enabled before you can use them, such as escalations, OpenNMS trap actions, and in vRealize Hyperic, SNMP notifications.

To define an alert action, select one of the tabs and supply the required information.

Escalation

Select an escalation from the Escalation Scheme pulldown; the tab refreshes and shows the escalation steps. You must define an escalation before you can assign it to an alert definition. Using an escalation that is configured to repeat until the alert is fixed is a good way to prevent redundant alerts firing for the same problem. To create an escalation, click **Escalation Schemes Configuration** on the **Administration** tab.

Control Action

In vRealize Hyperic, you can define a resource control action to perform when the alert fires. The control action can target the current resource (the one to which the alert definition is assigned) or a different resource on the same platform, as long as the resource type has supported control actions.

To configure a control action for the alert:

- 1 Select the **Control Action** tab and click **Edit**.

You can only assign a single control action to an alert definition.

You cannot assign a control action to a resource type alert.

Notify Roles

In vRealize Hyperic you can specify one or more roles as notification recipients. vRealize Hyperic users with a role you specify will be notified when an alert is fired.

- 1 Click **Add to List** on the **Notify Roles** tab.
- 2 On the roles selection page, choose the role(s) to be notified when the alert fires.

Notify other Recipients

Click **Add to List** on this tab to specify non-vRealize Hyperic user email recipients for alert notifications.

Script

In vRealize Hyperic, to assign a script action to the alert definition, click the **Script** tab, enter the full path to the script, and click **Set**. vRealize Hyperic will run the script when the alert fires. Scripts can reference alert-related vRealize Hyperic environment variables to perform custom notification logic.

OpenNMS

If vRealize Hyperic server is configured for OpenNMS integration, you can use this tab to configure vRealize Hyperic to send an SNMP trap to OpenNMS when the alert fires. The notification will be generated by `opennms_notify.gsp` alert notification template.

To configure an OpenNMS trap action, enter:

- Server - Listen address for the OpenNMS server.
- Port for the OpenNMS server.

SNMP Notification

If the vRealize Hyperic server is configured to send SNMP notifications to your NMS, you can use this tab to configure a trap notification action.

The notification sent when the alert fires will contain three variable bindings:

- `sysUptimeOID.0` - No configuration is required for this binding.
- `snmpTrapOID.0` - This binding is configured on the HQ Server settings page.
- A variable binding for the alert data specified in the `snmp_trap.gsp` alert notification template - the alert definition name and the short reason for firing.

To configure an SNMP notification action enter:

- IP Address - the address and port of the target NMS.
- OID -

The OID of the notification to send, which will contain the alert details specified in the `snmp_trap.gsp`, template.

- Notification Mechanism
 - v1 Trap
 - v2C Trap
 - Inform

Define a Recovery Alert for a Resource Type Alert

A recovery alert is special type of alert definition that you pair with a properly configured primary alert definition to streamline alert management.

The purpose of a recovery alert is to fire when the condition that fired another alert - the primary alert - is no longer true, and then mark the primary alert fixed and re-enable the primary alert definition. This strategy prevents redundant alerts and automates the task of marking an alert fixed.

You can define a recovery alert for a resource alert, and in vRealize Hyperic, a resource type alert. You cannot and don't need to define a recovery alert for a resource group alert in vRealize Hyperic - recovery alert behavior is automatic for resource group alerts.

To effectively leverage the benefits of recovery alert functionality you need to:

- Configure the primary alert definition to fire once when triggered and then disable itself until that fired alert is fixed. This prevents multiple alerts for a single incident.
- Configure a recovery alert definition and assign it to the primary alert definition. Make the recovery alert condition the opposite of the primary alert condition. The recovery alert fires when the primary alert condition is no longer true. Upon firing, the recovery alert marks the alert fired by the primary alert "fixed", and re-enables the primary alert definition, so that if the problem occurs again, the primary alert is again triggered.

Properly configured primary and recovery alert definitions keep users notified of problems without deluging them with alert notifications.

Create a Recovery Alert Definition for a Resource Type Alert

This topic provides instructions for creating a recovery alert definition for a resource type alert.

Procedure

- 1 Navigate to the resource for which the primary alert is assigned.
- 2 Click **New**.
- 3 Follow the directions in [Define a Resource Type Alert](#), setting the Condition Set to specify the condition that is the opposite of the primary alert definition's condition. For example if the primary alert condition is "1 Minute Load Avg > 2.0.", define the recovery alert condition as "1 Minute Load Avg < 2.0.
- 4 Use the **Recovery Alert** list to select the primary alert.

Using Dynamic Groups and Policies

6

You use dynamic groups and policies to manage resource monitoring in large scale environments.

- [Working with Dynamic Groups](#)

A dynamic group is a logical collection of resources that are organized by a specific criterion, such as name or type. After you have defined a dynamic group, you can assign one or more policies to manage large environments effectively.

- [Working with Policies](#)

A monitoring policy contains the settings and rules for monitoring resources of a specific type. You can assign a policy to one or more dynamic groups and you can edit a policy that has been assigned, or clone a policy as the basis of a new policy.

Working with Dynamic Groups

A dynamic group is a logical collection of resources that are organized by a specific criterion, such as name or type. After you have defined a dynamic group, you can assign one or more policies to manage large environments effectively.

After dynamic groups have been created, and monitoring policies defined and assigned, all the resources that match the group criteria become a member of that group and subject to the monitoring policies that are assigned to the group members. When vRealize Hyperic discovers a new resource, it is dynamically added to the group.

Note When you assign a policy to a dynamic group whose members are already configured, whether manually or by vRealize Hyperic auto-discovery, the policy becomes effective on the group members, meaning that configuration and monitoring parameters are overwritten with the policy configurations.

Using resource groups with assigned policies enables you to:

- Configure monitoring parameters, such as collected metrics and collection interval, for all group members.
- Manage access to group members according to user roles .
- Configure the resources that are members of the dynamic group.

Example: Adding a New Exchange Server to a Data Center

This scenario demonstrates how using dynamic groups works when a new Exchange server is added to a data center.

vRealize Hyperic discovers the Exchange server, after which you add it to the data center's inventory. vRealize Hyperic adds the server to the Exchange-type dynamic groups, and the server inherits the configurations and metrics that are specified in the Exchange monitoring policies.

Membership of Groups and Permissions

A resource is determined to be a member of a dynamic group when the following conditions exist:

- The resource matches the group selection criteria. For example, the resource name matches a regular expression.
- The resource is accessible in terms of permissions. One or more of the following must apply:
 - The resource is accessible by the owner of the dynamic group.
 - The resource is accessible by a user with a role that has been associated with the dynamic group. The role must have at least Write permissions for the groups.

Create a Dynamic Group

You can create groups and specify the criteria that defines when a resource is dynamically added to a group. You also assign the policies that determine the configuration and metrics for resource types, and roles that determine the access privileges for those resources.

You can create a dynamic group and choose to assign policies to it at a later time.

Prerequisites

- You must have sufficient privileges to create dynamic groups.
- If you are assigning roles, you must have sufficient privileges for this function.

Procedure

- 1 On the **Management > Dynamic Groups** tab, click the add icon.
- 2 Enter the necessary information in the Create Dynamic Group wizard.

Text boxes with red borders are mandatory. The **Next** button is not enabled until all mandatory information is provided.

- 3 On the Ready to Complete page of the wizard, review the criteria that you have defined.

This page is read only. If you must make changes, click **Back** to return to the relevant page.

The Members area shows the anticipated number of members in the group, based on the current inventory.

- 4 Click **Finish**.

On the **Resources** tab, the new dynamic group resource is added to the table. When you click the group name, a list of all the members of the group appears.

The effective policy that is applied to the resource appears on the **Inventory** tab of a resource that is a member of the group.

If you change the resource's configuration or the monitoring parameters after applying an policy, the policy is no longer applied on that resource. If the resource does not match the group criteria, it is excluded from the group and the policy is no longer effective. However, policy configurations continue to be applied on the resource, unless the resource configuration is modified.

Edit a Dynamic Group

You can edit definitions, for example the criteria or the policies that are assigned to a dynamic group.

When a group is updated, the policy is reapplied to all resources that are members of the updated group.

Prerequisites

- You must have sufficient privileges to edit dynamic groups.
- If you are assigning roles, you must have sufficient privileges for this function.

Procedure

- 1 On the **Management > Dynamic Groups** tab, select the group to change and click the edit icon.
- 2 In the Edit Dynamic Group wizard, use the **Next** button to navigate to the information that you are changing.
- 3 On the Ready to Complete page of the wizard, review the criteria summary.

This page is read only. If you must make additional changes, click **Back** to return to the relevant page.

- 4 Click **Finish**.

Delete a Dynamic Group

You can delete dynamic groups that are no longer required.

Prerequisites

You must have sufficient privileges to delete dynamic groups.

Procedure

- 1 On the **Management > Dynamic Groups** tab, select the group to remove and click the delete icon.
- 2 At the confirmation prompt, click **Yes**.

vRealize Hyperic removes the dynamic group. Resources that were members of the deleted group no longer have an effective policy. They retain their current configuration and do not receive any subsequent policy updates.

Working with Policies

A monitoring policy contains the settings and rules for monitoring resources of a specific type. You can assign a policy to one or more dynamic groups and you can edit a policy that has been assigned, or clone a policy as the basis of a new policy.

When you assign a policy to a dynamic group whose members are already configured, whether manually or by vRealize Hyperic auto-discovery, the policy becomes effective on the group members. An effective policy is a policy whose configuration and monitoring parameters were applied to the resource.

Configuration and Metrics

Settings can include configuration parameters and metric collection definitions.

Configuration Properties

Configuration parameters that are available for a resource type are inherited from the plug-in to which the resource is related. You can overwrite the default values when you create or edit a policy.

Some values are optional and others are mandatory.

Metric Collection

You can specify which of the available metrics of a resource type are to be collected, and the collection interval.

Default metrics are retrieved from **Administration > Monitoring Defaults**.

Effect of Changing Resource Configuration or Monitoring Parameters

If you manually change a resource's configuration or the monitoring parameters after applying a policy, the policy is no longer effective on that resource. However, if the policy or the dynamic group is updated later, the policy becomes effective on the resources again.

If a resource does not match the group criteria, it is excluded from the group and the policy is no longer effective. The configurations do not change unless the resource configuration is modified.

Policy Conflicts

In the event that there is a conflict between policies, for example when two dynamic groups containing different policies are relevant for the resource, the policy that is most recently updated, or the policy that was most recently assigned to a group, is the policy that applies. This process is followed even when a resource is manually changed.

Create a Policy

A policy is a configuration that can be assigned to a dynamic group.

Prerequisites

- You must have a role that grants full privileges to create policies.

- Roles that you are adding to a policy must have already been defined. See *Create a Role* in the *vRealize Hyperic Administration Guide*.

Procedure

- 1 On the **Management > Policies** tab, click the add icon.
- 2 Enter the necessary information in the Create Policy wizard.
Text boxes with red borders are mandatory. The **Next** button is not enabled until all mandatory information is provided.
- 3 On the Ready to Complete page of the wizard, review the values that you have defined.
This page is read only. If you must make changes, click **Back** to return to the relevant page.
- 4 Click **Finish**.

Edit a Policy

You can edit the settings and definitions of a policy.

When you edit a policy that has already been assigned to a group, the updated policy is reapplied to matching resources in that group.

Prerequisites

- You must have sufficient privileges to create policies.
- If you are assigning roles, you must have sufficient privileges for this function.

Procedure

- 1 On the **Management > Policies** tab, click the edit icon.
- 2 In the Edit Policy wizard, use the **Next** button to navigate to the information that you are changing.
- 3 On the Ready to Complete page of the wizard, review the criteria that you have defined.
This page is read only. If you must make changes, click **Back** to return to the relevant page.
- 4 Click **Finish**.

Clone a Policy

Cloning a policy enables you to use an existing policy as a template, adjusting the definition as required.

All parameters, other than the resource type can be changed when you clone a policy.

Prerequisites

- You must have sufficient privileges to create policies.
- If you are assigning roles, you must have sufficient privileges for this function.

Procedure

- 1 On the **Management > Policies** tab, click the clone icon.

- 2 In the Create Policy wizard, use the **Next** button to navigate to the information that you are changing.
- 3 On the Ready to Complete page of the wizard, review the criteria that you have defined.
This page is read only. If you must make changes, click **Back** to return to the relevant page.
- 4 Click **Finish**.

Delete a Policy

You can delete policies that are no longer required.

Prerequisites

You must have sufficient privileges to delete policies.

Procedure

- 1 On the **Management > Policies** tab, click the delete icon.
- 2 At the confirmation prompt, click **Yes**.

vRealize Hyperic removes the effective policy from the resources and from the policy list.

Global Monitoring and Alerting Settings



These sections describe options and settings that are defined globally for a vRealize Hyperic deployment.

- [Configuring Global Alert Properties](#)
This topic provides information about configuring global alert properties.
- [Configure Alert Notification Throttling](#)
You can use notification throttling to limit the number of alert email actions (notifications sent by email for a fired alert) that are issued in a 15 second interval.
- [Configuring Alert Notification Email Properties](#)
The settings in the Email Configuration Properties section of the Server Settings page are used to form notifications that vRealize Hyperic sends for a fired alert.
- [Configure Hyperic Version and Security Announcements](#)
vRealize Hyperic sends email announcements to vRealize Hyperic administrators when a key release is upcoming, or to distribute important product information.

Configuring Global Alert Properties

This topic provides information about configuring global alert properties.

The settings in the Global Alert Properties section of the **Administration > Server Settings** page enable immediate and global control of alert processing.

- Alerts - Disable or enable all alert definitions for all resources immediately. Disabling stops any alerts from firing; notifications defined in escalations that are currently in progress will be completed.
- Alert Notifications - Disable or enable alert notifications for all resources immediately. Disabling stops all notifications, include those for alerts with escalations currently in progress.
- Hierarchical Alerting - This setting controls whether alerts are evaluated using the hierarchical alerting method. When hierarchical alerting is enabled, before firing an alert for a resource, HQ considers the availability and alert status of the resource's parent. The purpose of hierarchical alerting is to avoid firing alerts for every resource affected by a single root cause.

Note You can extend the effect of hierarchical alerting by configuring the relationship between a network device or virtual host and the platforms that depend on it using the Network and Host Dependency Manager available in the Plugins section of the **Administration** tab.

Configure Alert Notification Throttling

You can use notification throttling to limit the number of alert email actions (notifications sent by email for a fired alert) that are issued in a 15 second interval.

When the threshold you specify is reached, vRealize Hyperic stops sending email alert notifications and instead sends a summary of alert activity every ten minutes to the recipients you specify.

After starting to throttle, vRealize Hyperic re-evaluates notification volume for fired alerts every 10 minutes; when it determines that the per interval volume of individual notifications that fired alerts would generate is less than the configured threshold, vRealize Hyperic resumes sending individual notifications.

In the Notification Throttling Configuration Properties section of the Server Settings page:

Procedure

- 1 Click the **Notification Throttling ON** control.
- 2 In the **Threshold** field, enter the maximum number of notifications you want sent in a 15 second interval.
- 3 Enter one or more email addresses in the **Notification Email(s)** field.

Configuring Alert Notification Email Properties

The settings in the Email Configuration Properties section of the Server Settings page are used to form notifications that vRealize Hyperic sends for a fired alert.

Table 7-1. Alert Notification Email Properties

Property	Description
Base URL	<p>The address:port where the vRealize Hyperic server listens for web application requests.</p> <p>The initial value of Base URL is the web application listen port configured when the vRealize Hyperic server was installed, for example:</p> <pre>http://Ms-MacBook-Pro-15.local:7080</pre> <p>Base URL forms the prefix of the URL to which Hyperic appends the remainder of the URL, which points to the Alert Detail page for the fired alert. For example:</p> <pre>http://Ms-MacBook-Pro-15.local:7080/alerts/Alerts.do?mode=viewAlert&eid=5:10611&a=16431</pre>
From Email Address	<p>The email address listed as the sender of the alert emails. For example:</p> <pre>hq@demo2.vmware.com</pre>

Configure Hyperic Version and Security Announcements

vRealize Hyperic sends email announcements to vRealize Hyperic administrators when a key release is upcoming, or to distribute important product information.

You can configure the level of messages you wish to receive or disable receipt of notifications with the Version and Security Announcements property.

Procedure

- ◆ In the Announcement Properties section of the Server Settings page, choose:
 - All
 - Major - Default value
 - None

Managing Alert and Notification Volume



The purpose of alerting is to speed the process of detecting and resolving problems. Rapid detection and response can be compromised when multiple alerts fire as a result of the same problem, or if responders are inundated by repetitive alert notifications.

Excessive alert and notification are less likely when:

- A given problem or root cause results in one, rather than many, alerts.
- An alert status of "unfixed" indicates a problem that still exists and needs attention, rather than a transient issue that occurred, and then went away.
- A single problem doesn't result in a firestorm of redundant notifications.

- [Preventing Multiple Alerts for the Same Problem](#)

When the volume of fired alerts is high, prioritizing and resolving problems is harder.

- [Disable all Alert Notifications](#)

If the volume of notifications exceeds manageable levels you can disable alert notifications globally.

- [Configure Network Host Dependencies for Hierarchical Alerting](#)

You can extend the reach of hierarchical alerting beyond the basic platform-server-service hierarchy to top level platforms - network devices or virtual hosts upon which operating system platforms depend.

- [Setting a Notification Throttle](#)

You can configure the vRealize Hyperic server to throttle back alert notifications in the event of an alert storm.

Preventing Multiple Alerts for the Same Problem

When the volume of fired alerts is high, prioritizing and resolving problems is harder.

You can reduce the overall volume of fired alerts without sacrificing visibility if you limit the number of times a given alert definition fires an alert for the same incident.

- Use repeating escalations - Assign an escalation that repeats until the alert is fixed. An alert in escalation cannot re-fire. The use of repeating escalations for all alerts is highly recommended and the best way to control alert volume in vRealize Hyperic.

- Fire one alert then disable the definition - You can configure a alert definition to fire once and disable itself until that alert is marked fixed. When the alert is marked "fixed" the alert definition is re-enabled. If you have vRealize Hyperic, you can can define an associated recovery alert to automatically fix the alert when the triggering condition is no longer true.

Disable all Alert Notifications

If the volume of notifications exceeds manageable levels you can disable alert notifications globally.

This option stops all alert notifications immediately including those resulting from escalations in process.

Procedure

- 1 Click the **Administration** tab.
- 2 Click **Server Settings**.
- 3 In the Global Alert Properties section, click the **Alert Notifications OFF** or **ON** control.

The change takes effect immediately. No alert notifications will be issued when OFF is selected. Escalations currently in progress will be terminated.

Configure Network Host Dependencies for Hierarchical Alerting

You can extend the reach of hierarchical alerting beyond the basic platform-server-service hierarchy to top level platforms - network devices or virtual hosts upon which operating system platforms depend.

To enable vRealize Hyperic to consider a top-level platform's availability and alert status before firing an alert for a dependent resources, you must define the relationship between a top-level platform and the operating system platforms that depend on it.

Procedure

- ◆ Use the Network Host Dependency Manager, available in the Plugins section on the **Administration** tab of the user interface.

Setting a Notification Throttle

You can configure the vRealize Hyperic server to throttle back alert notifications in the event of an alert storm.

You configure the maximum number of notifications that Hyperic will issue within a fifteen second interval. When the threshold is reached, vRealize Hyperic stops sending individual alert notifications, and instead, sends a summary of alert activity to designated recipients every ten minutes. When the volume of notifications falls below the specified threshold, vRealize Hyperic resumes sending individual notifications.

Notification throttling is disabled by default. You configure it on the Server Settings page, available from the **Administration** tab.

Manage Alert Definitions

This section provides information about accessing alert definitions, as well as enabling and disabling alert definitions.

- [Viewing and Editing Alert Definitions](#)

There are a variety of ways to navigate to pages that list and provide detailed information about alert definitions.

- [Enabling or Disabling Alert Definitions](#)

This section provides information about enabling or disabling alert definitions.

- [Enable Alert Definition Change Logging](#)

The vRealize Hyperic audit subsystem logs changes made to alert definitions. When an alert definition is enabled, disabled, or deleted, the event is written to the database. The information includes the username of the person who made the change.

Viewing and Editing Alert Definitions

There are a variety of ways to navigate to pages that list and provide detailed information about alert definitions.

- [List Alert Definitions for a Resource](#)

This topic provides instructions for listing alert definitions for a resource.

- [List Alert Definitions for a Resource Type](#)

This topic provides instructions for listing alert definitions for a resource type.

- [List Alert Definitions for a Resource Group](#)

This topic provides instructions for listing alert definitions for a resource group.

- [View All Alert Definitions for all Resources](#)

This topic provides instructions for viewing all alert definitions for all resources.

- [View and Edit an Alert Definition](#)

This topic provides instructions for viewing and editing an alert definition.

List Alert Definitions for a Resource

This topic provides instructions for listing alert definitions for a resource.

Procedure

- 1 Browse to the resource.
- 2 Click the **Alert** tab.
- 3 Click the **Configure** tab to see a list of alert definitions for the resource.

List Alert Definitions for a Resource Type

This topic provides instructions for listing alert definitions for a resource type.

Procedure

- 1 Click **Administration** in the masthead.
- 2 Click **Monitoring Defaults** in Server Settings section.
- 3 On the Monitoring Defaults Configuration page, click **Edit Alerts** for the resource type.

List Alert Definitions for a Resource Group

This topic provides instructions for listing alert definitions for a resource group.

Procedure

- 1 Browse to the resource group.
- 2 Click the **Alert** tab.
- 3 Click the **Configure** tab to see a list of alert definitions for the resource.

View All Alert Definitions for all Resources

This topic provides instructions for viewing all alert definitions for all resources.

Procedure

- 1 Click the **Analyze** tab in the Masthead.
- 2 Click **Alert Center**.
- 3 Click **Definitions**.

View and Edit an Alert Definition

This topic provides instructions for viewing and editing an alert definition.

Procedure

- ◆ Click the alert definition name in an alert definition list for the resource it is assigned to, or, click the alert definition name in the **Definitions** tab in the Alert Center.

Enabling or Disabling Alert Definitions

This section provides information about enabling or disabling alert definitions.

- [Enable or Disable a Resource Alert Definition](#)
There are multiple ways you can enable or disable a resource alert definition in the user interface.
- [Enable or Disable a Resource Type Alert Definition](#)
When you enable or disable a resource type alert, the change applies to all resources of that type.
- [Enable or Disable a Group Alert Definition](#)
This topic provides instructions for enabling or disabling group alert definitions.
- [Enable or Disable all Alert Definitions for Multiple Resources](#)
You can disable or enable all of the alert definitions for selected resources on the Browse Resource page.
- [Enable or Disable all Alert Definitions](#)
You can disable or enable alert definitions globally, if you want to turn alerting on or off for all resources in inventory.

Enable or Disable a Resource Alert Definition

There are multiple ways you can enable or disable a resource alert definition in the user interface.

Procedure

- 1 On the **Alerts** tab for a resource, browse the the resource to which the alert definition is assigned.
- 2 Click the **Alerts** tab.
- 3 Click **Configure** to display a list of alert definitions for the resource.

The page will include any alert definitions that the resource inherits from a resource type alert definition, marked with an asterisk. if you enable or disable the alert definition at this level, the change will be overridden by subsequent updates to the alert definition at the resource type level.

- 4 Select the checkbox for the alert definition you wish to enable or disable and toggle the **Set Active** control.

- 5 On the Alert Definition page:
 - a Navigate to the Alert Definition page by clicking the alert definition name on the Alert Center page.
 - b Click **Edit** in the Alert Properties section of the Alert Definition page.
 - c Toggle the **Set Active** button.

Enable or Disable a Resource Type Alert Definition

When you enable or disable a resource type alert, the change applies to all resources of that type.

Procedure

- 1 Click **Administration** in the masthead.
- 2 Click **Monitoring Defaults** in Server Settings section.
- 3 On the Monitoring Defaults Configuration page, click **Edit Alerts** for the resource type.
- 4 Click one or more alert definitions, and toggle the **Set Active** button.

Enable or Disable a Group Alert Definition

This topic provides instructions for enabling or disabling group alert definitions.

Procedure

- 1 Browse the the group to which the alert definition is assigned.
- 2 Click the **Alert** tab.
- 3 Click **Configure** to display a list of alert definitions for the group.
- 4 Select the checkbox for the alert definition you wish to enable or disable and toggle the **Set Active** control.

Enable or Disable all Alert Definitions for Multiple Resources

You can disable or enable all of the alert definitions for selected resources on the Browse Resource page.

Procedure

- 1 Click the **Resources** tab.
- 2 Use the filter options at the top of the page to list the resources of interest.
- 3 Select the checkbox for one or more resources.
- 4 Click **Enable all Alerts** or **Disable all Alerts**.

Enable or Disable all Alert Definitions

You can disable or enable alert definitions globally, if you want to turn alerting on or off for all resources in inventory.

Procedure

- 1 Click the **Administration** tab.
- 2 Click **Server Settings**.
- 3 In the Global Alert Properties section, click the **Alerts ON** or **OFF** control.

Enable Alert Definition Change Logging

The vRealize Hyperic audit subsystem logs changes made to alert definitions. When an alert definition is enabled, disabled, or deleted, the event is written to the database. The information includes the username of the person who made the change.

You can configure vRealize Hyperic server to also write the audit events to the `server.log` file in the server's logs directory.

Procedure

- ◆ Add the following section to `ServerHome/conf/server-log4j.xml`

```
<!-- Use this category to log every instance of a group alert firing.
      Comment out if the log messages become unwieldy. -->
<category name="org.hyperic.hq.galerts.processor.MemGalertDef.Fired">
  <priority value="DEBUG" />
</category>

<category
name="org.hyperic.hq.common.server.session.AuditManagerEJBImpl"> <!--SEE
HERE -->
  <priority value="DEBUG"/>
</category>
```

This section should be added before the root stanza.

The entries are written to `logs/server.log`, similar to the following:

```
2009-03-18 15:56:20,088 INFO
[main][org.hyperic.hq.common.server.session.AuditManagerEJBImpl@259]
  Audit Manager starting up2009-03-18 15:58:28,223 DEBUG
[main][org.hyperic.hq.common.server.session.AuditManagerEJBImpl@87]
Audit:Audit[user=admin,purpose=12288,time=1237409692209,resource=,
msg=HQ Started]2009-03-18 16:01:12,114 DEBUG
[UserLoginListener1][org.hyperic.hq.common.server.session.
AuditManagerEJBImpl@87]Audit:Audit[user=hqadmin,
purpose=16384,time=1237410072113,
resource=msg=HQ Administrator logged in]
```

```
2009-03-18 16:05:18,485 DEBUG [http-0.0.0.0-7080-4]
[org.hyperic.hq.common.server.session.AuditManagerEJBImpl@87]
Audit:Audit[user=hqadmin,purpose=20482,time=1237410318484,
resource=Linux,msg=Alert definition deleted (High Load)]
```


Creating and Managing Resource Groups and Applications

10

This section provides information about creating and managing resource groups and applications.

- [Configuring and Managing Resource Groups](#)
This section has instructions for creating and managing groups.
- [Defining Alerts for a Resource Group](#)
This section provides information about defining an alert for a compatible group - a group whose members are all of the same resource type.
- [Schedule Downtime for a Resource Group](#)
The Schedule Downtime page, available from the **Tools** menu when a resource, or a resource group is selected, allows you to schedule a downtime period, during which alerts for the resource(s).
- [Creating and Managing Applications](#)
This section provides information about configuring an application in vRealize Hyperic

Configuring and Managing Resource Groups

This section has instructions for creating and managing groups.

- [Create a Group](#)
This topic provides instructions for creating a group.
- [Add Resources to a Group](#)
This topic provides instructions for adding resources to a group.
- [Assign Roles to a Group](#)
This topic provides instructions for assigning roles to a group.
- [Remove Resources from a Group](#)
This topic provides instructions for removing a resource from a group.
- [Remove Roles from a Group](#)
This topic provides instructions for removing roles from a group.

Create a Group

This topic provides instructions for creating a group.

Procedure

- 1 From **Browse > Resource** page, click **New Group** on the **Tools** menu.
- 2 On the New Group page, enter the following:
 - Name
 - Description
 - Location
 - Make group private - select this checkbox to make the group private. A private group is invisible to other users, including admin users. You can share a private group with other users by associating it with a role. Note that the name you assign to a private group is automatically prefixed with the string private to *username*, where username is the creator's vRealize Hyperic username.
 - Contains Resources - Select the type of group:
 - Compatible/cluster - the group will contain resources of a single type. For example, Linux or JBoss 4.2.
 - Mixed - the group will contain multiple resource types. Mixed groups are useful for role-based resource access control.
 - The **Select Resource Type** pulldown appears to the right of the Contains Resources — the options vary depending upon whether you are created a compatible or mixed group.
 - For a compatible group, the **Select Resource Type** pulldown contains a list of all resource types in inventory.

For a mixed group, the **Select Resource Type** pulldown contains these choices:

 - Groups
 - Platforms, Servers, & Services
 - Applications
- 3 Click **OK**.

What to do next

Add resources to the new group. For more information, refer to [Add Resources to a Group](#).

Add Resources to a Group

This topic provides instructions for adding resources to a group.

Procedure

- 1 Navigate to the group's Inventory page, if it is not currently selected.
- 2 Click **Add To List** in the Resources section of the page.
- 3 On the Add to Group page, filter the resource list as desired by entering all or a part of the resource name in the **Filter By Name** field.
- 4 Select the checkbox for the desired resources and click the blue arrow to move them from the **Resources** column to the **Add Resources** column. (The arrow is enabled when you select a resource.)
- 5 Click **OK**.

Assign Roles to a Group

This topic provides instructions for assigning roles to a group.

Procedure

- 1 Navigate to the group's Inventory page, if it is not currently selected.
- 2 Click **Add To List** in the Roles section of the page.
- 3 On the Add to Roles to Group page, select the checkbox for the desired roles and click the blue arrow to move them from the **Roles** column to the **Add to Roles** column.
- 4 Click **OK**.

Remove Resources from a Group

This topic provides instructions for removing a resource from a group.

Procedure

- 1 Navigate to the group's Inventory page, if it is not currently selected.
- 2 In the Resources section, select the checkboxes for the resources you want to remove.
- 3 Click **Remove From List**.

Remove Roles from a Group

This topic provides instructions for removing roles from a group.

Procedure

- 1 Navigate to the group's Inventory page, if it is not currently selected.

- 2 In the Roles section, select the checkboxes for the roles you want to remove.
- 3 Click **Remove From List**.

Defining Alerts for a Resource Group

This section provides information about defining an alert for a compatible group - a group whose members are all of the same resource type.

- [Understanding Resource Group Alerts](#)

A resource group alert is an alert assigned to a compatible group - a group or resources you have defined that contains selected resources, all of which have the same resource type.

- [Define a Resource Group Alert](#)

This topic provides instructions for defining a resource group alert.

Understanding Resource Group Alerts

A resource group alert is an alert assigned to a compatible group - a group or resources you have defined that contains selected resources, all of which have the same resource type.

Resource group alerts are different than resource alerts in the following ways:

- A resource group alert definition has a single condition - A resource group alert definition is based on a single condition.
- A resource group alert condition is based on the behavior of multiple resources - Resource group alerts are evaluated differently than other alert types. A resource alert or resource type alert is fired for a specific resource based on monitoring results for that resource only. A resource group alert fires when a metric condition is true for a specified number or percentage of the resources in the group.
- The only action a resource group alert can trigger is an escalation - The only action you can assign to a group alert definition is an escalation.
- A resource group alert fixes itself - Although you can mark a group alert fixed - from the group's Alerts tab or other alert views - you don't need to. HQ automatically marks a group alert "fixed" when the condition that caused it to fire is no longer true. This shows up in the UI as "The problem fixed itself." in the "Alert Detail" screen.

Define a Resource Group Alert

This topic provides instructions for defining a resource group alert.

Procedure

- 1 Select a target compatible group.
 - a Browse to the compatible group to which the new alert definition will apply.
 - b Under the **Alert** tab, click **Configure**.
 - c Click **New** to display the New Alert page.

- 2 Define the alert properties.
 - **Name** — Name assigned by the user creating an alert definition. A fired alert is identified, in the user interface and alert notifications, by the alert definition name and a timestamp. An alert definition name should clearly communicate the nature of the problem. For example, Down for an alert on availability, or Low Memory for an alert on free memory.
 - **Description** — Description entered by the user creating the alert definition.
 - **Priority** — The severity of the problem, as defined by the person creating the alert definition: Low, Medium, or High. A consistent policy for defining an alert definition priority makes it easier to triage problems appropriately. An alert's priority is shown in vRealize Hyperic pages that present alert status and in alert notifications. You can sort alerts by priority in vRealize Hyperic's Alert Center or Operations Center.
 - **Active** — The current enabled/disabled status of the alert definition. Alerts only fire for enabled alert definitions. When an alert definition is disabled, Hyperic does not evaluate its condition or fire alerts for it.

- 3 Define the alert conditions that will cause the alert to fire.
 - a Select how many (an absolute number or percentage) of the resources in the group (the number of members is shown) that must satisfy the condition.
 - b Select the metric and the comparison value that those resources must satisfy.
 - c Click **OK**.

- 4 Assign the escalation for the alert.
 - a Click **Escalation Schemes Configuration** on the **Administration** tab to define an escalation.

The series of notifications defined in the escalation will be performed when an alert is fired. You must define an escalation before you can assign it to an alert definition.

Schedule Downtime for a Resource Group

The Schedule Downtime page, available from the **Tools** menu when a resource, or a resource group is selected, allows you to schedule a downtime period, during which alerts for the resource(s).

You define a period of time - a start and end date and time. At the start of the downtime period, currently active alert definitions for resources in the group are disabled. At the end of the period, those alert definitions are re-enabled.

Only a user with the Super User role may schedule downtime.

To schedule or reschedule maintenance:

Procedure

- 1 Navigate to the individual resource (a platform, server, or service) or the resource group for which you wish to schedule downtime.
- 2 Select **Schedule Downtime** from the **Tools** menu.
- 3 On the Schedule Downtime page, enter the start and end dates and times for the downtime and click **Schedule**.

Creating and Managing Applications

This section provides information about configuring an application in vRealize Hyperic

- [Understanding Instrument Java Components \(Optional\)](#)
You can gain deeper visibility into the health of Java applications by instrumenting application services.
- [Create an Application](#)
This topic provides instructions for creating an application.
- [Add Services to an Application](#)
This topic provides instructions for adding services to an application.
- [Add Application to a Group](#)
This topic provides instructions for adding an application to a group.
- [Monitoring an Application with Indicator Charts](#)
After you have created an application and added services to it, you can start monitoring all the application's components as a whole using indicator charts.

Understanding Instrument Java Components (Optional)

You can gain deeper visibility into the health of Java applications by instrumenting application services.

The vRealize Hyperic agent can auto-discover and manage Java application services via Model MBeans that adhere to a specified ObjectName naming convention and expose a specified set of service data. This enables deeper visibility into application health: you can monitor application services along with the hosting application server and its internal services.

Note Although instrumentation provides deeper visibility into Java application health, it is not required for application monitoring.

Create an Application

This topic provides instructions for creating an application.

Procedure

- 1 Click **Browse** on the **Resources** tab.
- 2 Click **New Application** in the **Tools** menu.
- 3 In the General Properties section of the New Application page, enter:
 - Name
 - Description
 - Location
- 4 In the Application Properties section of the New Application page, provide desired contact information:
 - Engineering Contact
 - Business Owner
 - IT Contact
- 5 Click **OK**.

Add Services to an Application

This topic provides instructions for adding services to an application.

Procedure

- 1 In the Services section of the Inventory page for the application, click **Add to List**.
The Edit Application page appears.
- 2 Select the checkbox for the desired services and click the blue arrow to move them from the **Services** column to the **Add Services** column. (The arrow is enabled when you select a resource.)

- 3 Click **OK**.

Add Application to a Group

This topic provides instructions for adding an application to a group.

Procedure

- 1 In the Groups section of the Inventory page for the application, click **Add to List**.
The Edit Application page lists existing mixed groups of applications.
- 2 Select the checkbox for the desired services and click the blue arrow to move them from the **Services** column to the **Add Services** column. (The arrow is enabled when you select a resource.)

Monitoring an Application with Indicator Charts

After you have created an application and added services to it, you can start monitoring all the application's components as a whole using indicator charts.

Initially, the Monitor tab displays only the availability metric indicator for the new application. The services you added are listed, along with Host Servers (the servers that host the listed services), on the Resources minitab tab, and through those you can add more metrics to the display.

Map vCenter Virtual Machine IDs to vRealize Hyperic Platforms

11

In order for vRealize Operations Manager to synchronize information from a vCenter server with vRealize Hyperic, the vRealize Hyperic platforms must have a unique identifier assigned that can be mapped to the corresponding vCenter virtual machine.

To configure unique IDs for vRealize Hyperic platforms:

Procedure

- 1 On the masthead, click the **Administration** tab.
- 2 Select the **HQ Server Settings** link.
- 3 In the vCenter Server Settings section, enter the appropriate values:
 - a Enter the URL of the vCenter SDK in the **vCenter URL** text box.
The URL is generally `https://<IP address of vCenter Server>/sdk`
 - b Enter the username of a user with vCenter administrator privileges in the **vCenter User** text box.
 - c Enter the password for the specified user in the **vCenter Password** text box.
- 4 Click **OK**.
- 5 (Optional) Verify that the properties have been correctly applied.
 - a On the **Resources** tab, click a platform link.
 - b Verify that vCenter UUID and MOID values appear in the summary information at the top of the page.

Configuring Hyperic as an SNMP Trap Receiver

12

You can set up vRealize Hyperic to receive and log SNMP messages from a remote network management system or other SNMP-capable system or device.

You can set alerts based on the received SNMP data to trigger notifications or other responses.

The process involves configuring the agent and the network system or device to communicate with each other, and creating a platform of type Network Device to represent the remote system or device.

- [Step 1 - Select and Enable an Agent to Receive Traps](#)

When you configure a platform of type Network Device, you must specify connection information for a vRealize Hyperic agent running on an operating system in your environment - that agent will receive the SNMP data. The agent you designate must be able to receive traps.

- [Step 2 - Configure Trap Generator](#)

This topic provides instructions for configuring the trap generator.

- [Step 3 - Create a Proxy Platform for Trap Data](#)

In vRealize Hyperic, remotely monitored resources are represented as a platform.

- [Resolving Configuration Issues](#)

This topic provides information about dealing with configuration issues.

Step 1 - Select and Enable an Agent to Receive Traps

When you configure a platform of type Network Device, you must specify connection information for a vRealize Hyperic agent running on an operating system in your environment - that agent will receive the SNMP data. The agent you designate must be able to receive traps.

The agent's default UDP port for receiving traps is 162, in the privileged range, which the agent cannot access unless it runs as root (or as an Administrative user on Windows).

To configure a non-privileged port to receive SNMP traps:

Procedure

- 1 Open the `agent.properties` file in a text editor.

- 2 Add the following line to the file:

```
snmpTrapReceiver.listenAddress=udp:0.0.0.0/1620
```

- 3 Restart the agent.

This setting enables the agent to receive traps through any interface on the platform at UDP port 1620. If desired, you can specify an specific interface's IP address, or different port in the unprivileged range (1024 or above).

If you try to check that the SNMP port is open at this point, netstat will not show it open. Once you complete all steps in this procedure netstat output will produce expected results.

Step 2 - Configure Trap Generator

This topic provides instructions for configuring the trap generator.

Procedure

- 1 Set up the agent as a trap destination on the SNMP-enabled system or application.
- 2 Supply the agent's SNMP listen address and port - either the default connection settings, or if applicable, those configured in the previous step.

Step 3 - Create a Proxy Platform for Trap Data

In vRealize Hyperic, remotely monitored resources are represented as a platform.

To create a platform for the remote device or host:

Procedure

- 1 Select **New Platform** from the **Tools** menu.
- 2 Enter a name for the platform.
- 3 Select **Network Device** from the **Platform Type** pulldown.
- 4 Enter the FQDN of the network device or host.
- 5 Select the agent connection from the pulldown list - it must have connectivity to the SNMP port on the device or host.

- 6 Enter the IP address of the network device or host.

Do not use 127.0.0.1, even if the application is local to the agent being used for the SNMP connection.

- 7 The netmask or MAC address of the device or server may be required depending on the configuration of the device.
- 8 Click **OK**.

The Inventory tab for the new platform is displayed.

- 9 Click **Edit** in the Configuration Properties section.
- 10 Select the interface.index the device uses to uniquely identify interface services. Most devices work with the default (ifDescr).
- 11 Supply or validate the snmplp, snmpPort and SNMP community string.
- 12 Select the value for the snmpVersion that corresponds to the device.
- 13 If you use SNMP v3, it is likely that authentication is configured, in which case enter username, password and auth type.
- 14 Select the **Enable Log Tracking** checkbox.
- 15 Click **OK**.

Resolving Configuration Issues

This topic provides information about dealing with configuration issues.

Problem

SNMP configuration is not accepted

Solution

- 1 Verify the IP address, SNMP port, SNMP version, and authentication settings, are properly configured.
- 2 Verify that the agent you selected for the proxy platform has network connectivity and can reach the SNMP port of the network device (check firewalls).
- 3 Make sure that the community string configured for the proxy platform is the same as the community string the trap generator uses to send traps. If the strings do not match, traps cannot be sent to the proxy platform - instead, traps will be routed to the default platform, if that platform has log tracking enabled.

Check the vRealize Hyperic agent log file to make sure that the agent has opened the proper UDP port. If properly configured, you should see messages similar to:

```
2008-06-05 16:45:05,447 DEBUG [SNMPTrapReceiver] snmpTrapReceiver.listenAddress=udp:0.0.0.0/1620
2008-06-05 16:45:05,572 DEBUG [SNMPTrapReceiver] Add 1:10003 for 10.2.0.2-switch
```

Managing Hyperic Users and Roles

13

This section provides information about managing users and roles.

- [Creating and Manage User Accounts](#)
This topic provides information about managing users and accounts.
- [Creating and Managing Roles in vRealize Hyperic](#)
This section provides information about creating and managing roles.
- [Understanding Role-based Dashboards](#)
A role-based dashboard is a version of the vRealize Hyperic dashboard available to users assigned to the role.

Creating and Manage User Accounts

This topic provides information about managing users and accounts.

If vRealize Hyperic is configured to use your enterprise directory for user authentication, you do not have to manually create user accounts. vRealize Hyperic will automatically create an account for users that are externally authenticated by your LDAP, Active Directory, or Kerberos system.

Create a User Account

This topic provides instructions for creating a user, and assigning roles to the user.

Procedure

- 1 Click **New User** on the Administration page.
- 2 Enter values for the following fields:
 - Name
 - Username
 - Phone
 - Department
 - Password - Passwords must contain at least 6 case-sensitive characters and numbers, and no spaces or quotation marks.

- Email
- Format - can be either HTML or plain text.
- SMS Address — An email-to-SMS gateway email address for the user's SMS device.

For a cellular phone on the Cingular network, this might look like 4155551212@mobile.mycingular.com. Check with the service provider for details about an email-to-SMS configuration. Basic alert notification sent to this user's SSS address will be in long format, which can result in up to five separate messages on the SMS device each time notification is sent by vRealize Hyperic. vRealize Hyperic recommends that SMS alerting be used in conjunction with escalations, not basic alert notification, but short format is used there.

- Enable Login - Toggle the radio button to disable or enable the account. The user cannot log in when the login is disabled

3 Click **OK and Assign to Roles**.

Assign Roles to a New User Account

This topic provides instructions for assigning roles to a new user account.

In vRealize Hyperic, the roles to which a user is assigned govern which resources the user may access, and the operations the user can perform on those resources. Each role in vRealize Hyperic defines a permission matrix; users with a role may exercise the permissions it grants, on the resources in groups assigned to the role.

Procedure

- 1 In the Roles panel on the left side of the page, select the checkbox for each role to which you want to assign the user.
- 2 Click the blue arrow to move the roles to the Add to Roles panel.
- 3 Click **OK**.

View User Accounts

This topic provides instructions for listing user accounts.

Procedure

- 1 Click **List Users** on the Administration page.

The List Users page lists the following information for each user account.

- First Name
- Last Name
- Username
- Email
- Department

- 2 Click the Username of the users whose details you want to view.

The following information is displayed:

- Name
- Username
- Phone
- Department
- Password - The user's password is not displayed. If you have the permission to modify the password, a **Change** link is presented.
- Email
- Format - Format for email notifications sent to the user - HTML or plain text.
- SMS - An email-to-SMS gateway email address for the user's SMS device.
- Enable Login — Indicates whether or not the account is enabled. The user cannot log in when the login is disabled.
- Roles assigned to - this section contains a list of roles to which the user is assigned.

- 1 To assign additional roles, click **Add to List**.

Change a User Password

In vRealize Hyperic, only a user with a role that grants Full or Read/Write permissions to user accounts may change another user's password.

You must enter your current password to change an account password, whether the account is yours, or another user's.

Procedure

- 1 Click the **Administration** tab in the user interface.
- 2 Navigate to the user account for the user whose password you want to change. For more information, see [View User Accounts](#).
- 3 In the General Properties section of the page that appears, click the **Change...** link in the Password field.

The Change Password page appears.

- 4 Enter your password in the **Enter Your Current Password** field.

The password entered must be the password of the user making the change, not the password of the user whose password is being changed.

- 5 Enter the new password for the selected account in the **New Password** field.
- 6 Re-enter the password in the **Confirm New Password** field.

- 7 Click **OK**.

Edit Account Settings

This topic provides information about editing account settings.

Procedure

- 1 Click **Edit** in the General Properties section of the UserName page.
- 2 Edit the respective fields.
- 3 Click **OK**.

Update Roles Assigned to a User

This topic provides instructions for updating roles assigned to a user.

Procedure

- 1 Click **Add to List** in the Roles Assigned To section of the UserName page. The Edit UserName page appears.
- 2 In the Roles panel on the left side of the page, select the checkbox for each role to which you want to assign the user.

Click the blue arrow to move the roles to the Add to Roles panel.
- 3 Click **OK**.

Creating and Managing Roles in vRealize Hyperic

This section provides information about creating and managing roles.

If you use LDAP or Active Directory for user authorization, vRealize Hyperic can automatically assign users to vRealize Hyperic roles whose name matches a group to which the user is assigned - if the role names and LDAP group names adhere to a required naming convention. To take advantage of this functionality:

- a role name in vRealize Hyperic must match a group name in your authorization directory. In addition to being identical, both role name and group name must be prefixed by org\, for example org\Admin.
- vRealize Hyperic server must be configured to use your LDAP system.

Create a Role

This topic provides instructions for creating a new role.

Procedure

- 1 Click **New Role** on the Administration page.
- 2 In the Properties section of the New Role page, enter the name and description for the role.

3 In the Permissions section, select a permission level:

- **Full**
- **Read-Write**
- **Read-Only**
- **None**

For each type:

- Users
 - Grant **Full** to enable role users to create and delete user accounts.
 - Grant **Read-Write** to enable role users to edit users accounts.

- Roles

If you select **Full**, which enables role users to create roles, vRealize Hyperic will ensure that the role's permission level to Users and Groups is at least **Read-Only**, because to create a role, you need to view users and groups.

- Groups

- Grant **Full** to enable role users to delete groups created by other users.
- Grant **Read-Write** to enable role users to modify groups created by others.

Regardless of the permission level you select, any user can create groups, and as the owner of such groups, delete them.

- Platforms

- If you select **Full**, which enables role users to delete platforms and their child resources, vRealize Hyperic will require that the role's permission level to Servers and Services is also **Full**.
- If you select **Full** or **Read-Write**, vRealize Hyperic will automatically select the checkbox for the Can Fix/Ack Alerts? and Can Control? capabilities.
- If you select **Read-Only**, you have the option to grant alert management or resource control capabilities by clicking Can Fix/Ack Alerts? or Can Control? respectively.
- If you select **None**, you cannot grant alert management or resource control permissions.

- Servers

- If you select **Full**, which enables role users to delete servers and child services, vRealize Hyperic will require that the role's permission level to Platforms is at least **Read-Write**, and its permission level to Services is **Full**.
- If you select **Full** or **Read-Write**, vRealize Hyperic will automatically checkmark the Can Fix/Ack Alerts? and Can Control? capabilities.
- If you select **Read-Only**, you have the option to grant alert management or resource control capabilities by clicking Can Fix/Ack Alerts? or Can Control? respectively.

- If you select **None**, you cannot grant alert management or resource control permissions.
- Services
 - If you select **Full**, vRealize Hyperic will require that the role's permission level to Servers is at least **Read-Write**.
 - Grant at least **Read-Only** if you are going to grant the role Full permission to Applications.
 - If you select **Full** or **Read-Write**, vRealize Hyperic will automatically select the checkboxes for the Can Fix/Ack Alerts? and Can Control? capabilities.
 - If you select **Read-Only**, you have the option to grant alert management or resource control capabilities by clicking Can Fix/Ack Alerts? or Can Control? respectively.
 - If you select **None**, you cannot grant alert management or resource control permissions.
- Applications
 - Grant **Full** if you want role users to be able to create and delete applications.
 - Grant **Read-Write** if you want role users to be able to modify change applications created by others.
- Escalations
 - Grant **Full** if you want role users to be able to create and delete escalations groups.
 - Grant **Read-Write** if you want role users to be able to modify escalations.
- Policies
 - Grant **Full** if you want role users to be able to create and delete policies.
 - Grant **Read-Write** if you want role users to be able to modify policies created by others.
 - Grant **Read-Only** if you want role users to be able to view policies created by others.
 - If you select **None**, you cannot grant policies permission.

The role is saved, and the refreshed role page will have three new sections: Assigned Users, Assigned Groups, and Alert Calendar.

What to do next

Once you have saved the role in the system, proceed to [Assign Users to a Role](#).

Assign Users to a Role

In vRealize Hyperic, each user you assign to a role can exercise the permissions defined for the role, on resources in the groups assigned to the role.

Procedure

- 1 If you are not currently viewing the role to which you wish to assign users, navigate to the role.
- 2 Click **Add to List** in the Assigned Users section of the page.

The Assign Users to Role page appears.

- 3 On the Users panel on the left side page, select the checkbox for each user you wish to add to the role, and click the blue arrow to move the users to the Assign To Role panel.
- 4 Click **OK**.

What to do next

If you are creating a role purely for the purpose of role-based alert notification, skip to [Define Alert Calendar for Follow-the-Sun Role-Based Notifications](#). Otherwise, proceed to [Assign Groups to a Role](#).

Assign Groups to a Role

In vRealize Hyperic, the groups of resources you assign to a role constitute the population of resources to which users with that role may exercise the permissions associated with the role.

For example, if the role's permission level to Platforms is None, role users will not have access to platforms in groups assigned to the role.

Assigning a group of groups or a group of applications to a role does not grant permissions to individual resources in nested groups or applications. Groups you assign to a role for the purpose of granting permissions must contain individual resources, as opposed to other groups or applications.

Procedure

- 1 If you are not currently viewing the role to which you wish to assign groups, navigate to the role.
- 2 Click **Add to List** in the Assigned Groups section.
The Assign Groups to Role page appears.
- 3 On the Groups panel on the left side of the page, select the checkbox for each resource group you wish to add to the role, and click the blue arrow to move the groups to the Assign To Role panel.
- 4 Click **OK**.

Define Alert Calendar for Follow-the-Sun Role-Based Notifications

An alert calendar defines the availability calendar during which role users are available for alert notifications.

You should define an alert calendar if:

- You are creating a role that will be a recipient of alert notifications, and
- The users assigned to the role users are available only during specific intervals only.

By default, a role's alert calendar settings specify that role users are available for notifications 24 hours a day, 7 days a week, with no exceptions. To define a narrower availability calendar:

Procedure

- 1 For each day in the week,
 - Use the first set of From and To pull-downs to specify a start time and an end time that role users are availability for notifications.

- If there is a period of time within the availability period specified in the previous step, during which role users should not receive notifications, click **Except**, and use the **From** and **To** pull-downs on the right to specify that period of time.

2 Click **Save**.

You must define additional role or roles with complementary alert calendars to ensure that there is a role whose users are available during periods of time that the current role's alert calendar does not include.

Customize Role-Specific Dashboard

When you create a role, vRealize Hyperic creates a Dashboard with the same name as the role.

Procedure

- ◆ Select the customized dashboard from the **Select a Dashboard** pull-down in the upper left corner of the vRealize Hyperic dashboard.

Use Roles for Follow the Sun Alerting

vRealize Hyperic allows you to notify users of incidents based on support schedules, greatly simplifying the management of 24x7 or rotating support teams.

This is accomplished through the use of role-based alert calendars. The alert calendar can be most usefully applied in an , wherein multiple sets of people can be selected for notification, but only the one whose alert calendar is currently open will be notified.

Procedure

- 1 Create multiple roles.
- 2 For each role, assign complementary alert calendars (that is, alert calendars that span different periods of time, but all together cover the calendar). On the Edit Role screen, follow the instructions for creating a calendar for a role.
- 3 Define an escalation scheme with at least one SMS or email escalation action.
In the escalation action, instruct vRealize Hyperic to notify All users assigned to a specific role and then select all of the just-defined roles.
- 4 Create an alert and assign that escalation scheme to it.

When the alert is fired, vRealize Hyperic will start performing the escalation actions but will only notify the roles whose alert calendars are currently open.

Role Definitions

This topic provides information about the data available when viewing a role.

Properties

The Properties section contains this information:

- Name
- Owner
- Description
- Administer Server Configuration
- Dashboard Name

Permissions

The Permissions section displays the permission matrix for the role.

- • None - No access at all to instances of the type.
- • Read-Only - Allows role users to view instances of the type, but not create, edit, or delete them.

For Platforms, Servers, Services, Groups, also enables:

- Read-Only access to alert definitions for the inventory type.

A role with Read-Only permission level does not have permissions to enable/disable/fix/ack alerts or control resources - these capabilities must be explicitly granted.

- Read-Write - Allows role users to view and edit instances of the type, but not create or delete them. For Platforms, Servers, Services, Groups, also gives:
 - Full access to alert definitions for the inventory type
 - Permission to manage alerts (enable/disable, fix, acknowledge) for the inventory type.
 - Permission to perform supported control operations on resources of the inventory type.
- • Full - Allows role users to create, edit, delete, and view instance of the type.

For Platforms, Servers, Services, Groups, also gives:

- Full access to alert definitions for the inventory type.
- Permission to manage alerts (enable/disable, fix, acknowledge) for the inventory type.
- Permission to perform supported control operations on resources of the inventory type.

Assigned Users

The Assigned User section lists the following information for each user assigned to the role:

- First Name

- Last Name
- Username - you can click the username to view the user.

Assigned Groups

The Assigned Groups section lists the group name and description of the group for each group assigned to the role.

Alert Calendar

The Alert Calendar section defines the availability calendar during which role users are available for alert notifications.

Custom alerts calendars enable follow-the-sun alert notifications. To implement a follow the sun notification strategy, you create multiple roles with complementary alert calendars. You assign a user to the role whose alert calendar matches the user's availability.

By default, a role's alert calendar settings specify that role users are available for notifications 24 hours a day, 7 days a week, with no exceptions.

To define an availability calendar for a role:

- 1 For each day in the week,
 - a Use the first set of From and To pull-downs to specify a start time and an end time that role users are availability for notifications.
 - b If there is a period of time within the availability period specified in the previous step, during which role users should not receive notifications, click **Except**, and use the From and To pull-downs on the right to specify that period of time.
- 2 Click **Save**.

You must define additional role or roles with complementary alert calendars to ensure that there is a role whose users are available during periods of time that the current role's alert calendar does not include.

Understanding Role-based Dashboards

A role-based dashboard is a version of the vRealize Hyperic dashboard available to users assigned to the role.

A role-based dashboard can be customized to match the needs and interests of users with that role. You can configure the contents of a role-based dashboard to meet the needs of users with that role.

A role-based dashboard is available in addition to a user's personal dashboard. Every vRealize Hyperic user has a personal dashboard, and a role-based dashboard for each role to which he is assigned.

Any modification to a role-based dashboard affects all users with the role; unlike a personal dashboard, a role-based dashboard cannot be customized on a per user basis.

Role-Based Dashboards are Automatically Available

Upon creation of a new role, a copy of the standard vRealize Hyperic dashboard is saved and available to members of the role. Any user with this role can navigate to it using the Select a Dashboard pull-down, and designate it as his default dashboard.

The default dashboard for a new user is his personal dashboard. To set a role-based dashboard to be the default, a user must select it and click **Make Default**.

Permissions for Modifying Role-Based Dashboards

A role-based dashboard can be modified by:

- an administrator with the Super User role.
- a user assigned to the role, if the role grants the Modify permission.

Tailor a Role-Based Dashboard

This topic provides information about customizing a role-based dashboard.

To tailor the layout and content of a role-based dashboard:

Procedure

- 1 Select the role-based dashboard from the **Select a Dashboard** pull-down on the currently displayed dashboard.
- 2 Add, configure, remove, and rearrange the portlets in the dashboard, as desired.

Modifications to a role-based dashboard are immediate, and will appear any time a user with that role accesses it.

View and Manage Resource Data

14

In the vRealize Hyperic user interface, a resource's Inventory Page presents a variety of resource properties and key information about the resource's closest relatives in the resource hierarchy.

The information on the Inventory page varies by both inventory type and resource type, and includes auto-discovered and user-configured resource properties. For a fundamental inventory type — a platform, service, or application — the Inventory Page contains a Configuration Properties section, where you can view and update the properties that enable or tailor monitoring behaviors for the resource.

- [Navigate to the Inventory Page for a Platform](#)

This topic provides instructions for accessing the Inventory page for a platform.

- [Viewing Inventory Properties for a Platform](#)

This topic describes the data and controls that appear at the top of the Inventory page for a platform.

- [Navigate to the Inventory Page for a Server](#)

This topic provides instructions for accessing the Inventory page for a server

- [Viewing Inventory Properties for a Server](#)

This section describes the contents of the Inventory page for a server.

- [Navigate to the Inventory Page for a Service](#)

This topic provides instructions for accessing the Inventory page for a service

- [Viewing Inventory Properties for a Service](#)

This topic describes the data and controls that appear at the top of the Inventory page for a service.

- [Navigate to the Inventory Page for a Group](#)

This topic provides instructions for accessing the Inventory page for a group.

- [Viewing Inventory Properties for a Group](#)

This topic describes the data and controls that appear at the top of the Inventory page for a group.

- [Navigate to the Inventory Page for an Application](#)

In vRealize Hyperic, an application is an inventory type, configured by an authorized user. An application is a set of selected services, usually running in different servers on multiple platforms, that together fulfill a single business purpose.

- [Viewing Inventory Properties for an Application](#)

This topic describes the data and controls that appear at the top of the Inventory page for an application.

- [Creating Network Entities Manually](#)

Most resource types are automatically discovered, and auto-discovery is the only way an operating system platform can be added to inventory. Conversely, auto-discovery of network or virtual platform types is not possible — you must add such platforms to inventory manually.

- [Remove a Resource from the Inventory](#)

This topic provides instructions for removing a resource from the vRealize Hyperic inventory.

Navigate to the Inventory Page for a Platform

This topic provides instructions for accessing the Inventory page for a platform.

Procedure

- 1 Click **Resources > Browse > Platforms** to navigate to the platform.
- 2 Click the **Inventory** tab.

Viewing Inventory Properties for a Platform

This topic describes the data and controls that appear at the top of the Inventory page for a platform.

Resource Properties

The properties at the top of the Inventory page for a platform provide identifying information about the managed product. The inventory properties displayed for a platform vary slightly by platform type, but typically will include most of these properties.

The properties displayed are:

- Description
- Owner - By default, the vRealize Hyperic user under whose account the resource was added to inventory. Click **Change...** to assign a different resource owner.
- Secondary DNS
- Default Gateway
- Vendor
- Vendor Version
- IP address
- Primary DNS
- CPU Speed

- OS Version
- RAM
- Architecture

The plugin developer controls which resource properties are displayed at the top of a resource's Inventory page: any properties enclosed in a <properties> element for a resource type appear in the page header when you browse to an instance of that type.

Map Control for a Platform

Click the Map control in the page header for a graphical view of the servers on the platform. The platform is displayed at the bottom of the map; servers are displayed above the platform. Resource names are in bold and the resource type is displayed just below in smaller type. Click a resource name to navigate to that server.

Tools Menu for a Platform

When a platform is selected, the Tools menu has the following commands:

- Configure Platform — Opens the Configuration Properties page for the platform, where you can edit the resource's configuration properties.
- Clone Platform — Creates a new platform with the same configuration as an existing platform. For more information, refer to [GUID-3C091A0B-340C-4D02-8F92-10BAA6C05C09#GUID-3C091A0B-340C-4D02-8F92-10BAA6C05C09](#).
- Delete Platform — Delete the platform, its platform services, servers, and the services in the servers from inventory.

If you delete all of the platforms that a vRealize Hyperic agent manages, the vRealize Hyperic server also removes the saved authentication token for that agent from the vRealize Hyperic database. Therefore, if you delete a platform that is managed by an agent that does not manage any other platforms (as in the case of an agent that manages only the platform it runs on), the agent will no longer be able to connect to the server. If you want the agent to rediscover the platform where the agent runs, you must repeat the agent setup process. Otherwise (if you do not want the agent to re-discover the platform it runs on) shut down the agent, and uninstall the agent from that platform.

- New Server — Displays the New Server page, where you can manually add a new server to the platform. (For instance, a server that was not auto-discovered.) For more information, refer to [Create a Server](#).
- New Platform Service — Displays the New Service page, where you can manually add a new service, for instance a remotely monitored network service, to the platform. For more information, refer to [Create a Platform](#).
- New Auto-Discovery
- Enable All Alerts on this Agent

- Disable All Alerts on this Agent
- Add to Dashboard Favorites — Adds the platform to the Favorite Resources portlet on the current user's dashboard.
- Add to Group — Opens the Group Manager page, which lists the groups to which the platform may be added. You can add a platform to a group if:
 - You have permission to access the group, and:
 - The group is a mixed group that contains platforms, servers, and services.
 - The group is a compatible group of the selected platform's platform type.

General Properties for a Platform

The General Properties section of the Inventory page for a platform lists the following information:

- Description
- Date Created
- Location
- Date Modified
- Resource Type

Type and Network Properties for a Platform

The Type and Network Properties section of the Inventory page for a platform lists the following information:

- Platform Type
- Agent Connection — The IP address:port pair that the vRealize Hyperic server will use to connect to the agent on the platform device.

In the case of a platform type that is monitored by an agent on a different platform, such as a network or virtual host, this property identifies the agent that manages the platform.

- Fully Qualified Domain Name — The platform's FQDN.
- IP Address, MAC Address, Netmask — One or more sets of these identifiers for the platform. There is at a minimum one set for the loopback (local) IP address — 127.0.0.1 — and then additional sets for each network interface on the device.

Servers on a Platform

The Servers section lists the following information for each server on the platform:

- Server
- Server Type
- Install Path

- Description
- Availability

There are two controls available:

- New — Click to add a server to the platform.
- Delete — Click to delete a server from the platform. Services in the server will also be deleted.

Services on a Platform

This section lists the following information for each platform service on the platform:

- Service - The resource name, presented as a hyperlink you can click to navigate to the service.
- Description
- Availability

There are three controls available:

- View — This pull-down allows you to filter the list to display only platform services of a selected type.
- New — Click to add a platform service to the platform.
- Delete — Click to delete a platform service from the platform.

Groups Containing a Platform

This section lists the following information for each the group of which the platform is a member:

- Group — The resource name, presented as a hyperlink you can click to navigate to the group.
- Description

There are two controls available:

- Add to List — Click to add the platform to a group.
- Remove From List — If you have checkmarked one or more groups in the list, this control allows you to remove the resource from the selected groups.

Configuration Properties for a Platform

This section displays the configuration properties for the platform:

- Shared — These properties are typically only present for platform types that vRealize Hyperic does not auto-discover, and vary by platform type.
- Monitoring — These properties control log and configuration tracking for the platform.

Navigate to the Inventory Page for a Server

This topic provides instructions for accessing the Inventory page for a server

Procedure

- 1 Click **Resources > Browse > Servers** to navigate to the platform.
- 2 Click the **Inventory** tab.

Viewing Inventory Properties for a Server

This section describes the contents of the Inventory page for a server.

Resource Properties

The properties at the top of the Inventory page for a server provide identifying information about the managed product. The inventory properties displayed for a server vary by server type, but typically will include vendor name and software version.

The properties displayed in the screenshot below (for a JBoss server) are:

- Java Vendor
- Owner - By default, the vRealize Hyperic user under whose account the resource was added to inventory. Click **Change...** to assign a different resource owner.
- Build Date
- JBoss Version
- Version Name

The plugin developer controls which resource properties are displayed at the top of a resource's Inventory page: any properties enclosed in a <properties> element for a resource type appear in the page header when you browse to an instance of that type.

Map Control for a Server

Click the Map control in the page header to view the server's child services(s) and its parent platform.

You can click the name of a child or parent resource to view its Monitor page.

Tools Menu for a Server

When a server is selected, the **Tools** menu has the following commands:

- **Configure Server** — Opens the Configuration Properties page for the server, where you can edit the resource's configuration properties.
- **Delete Server** — Delete the server and all its child services from inventory.

- **New Service** — Displays the New Service page, where you can manually add a new child service to the server. (For instance, a service that was not auto-discovered.)
- **Add to Dashboard Favorites** — Adds the server to the Favorite Resources portlet on the the current user's dashboard.
- **Add to Group** — Opens the Group Manager page, which lists any compatible groups of the same type as the server.

General Properties for a Platform

The General Properties section of the Inventory page for a platform lists the following information:

- Description
- Date Created
- Date Modified
- Resource Type

Type and Host Properties

The Type and Host Properties section of the Inventory page for a server lists the following information about the server:

- Install Path
- Host Platform

Services on a Server

The Services section of the Inventory page for a server contains:

- The total number of child services on the server, and the number of services of each service type discovered on the server.
- The following columns for each service:
 - Service
 - Service Type
 - Description
 - Availability

Groups Containing a Server

The Groups Containing This Resource section of the Inventory page for a server lists the groups of which this server is a member. To view the Inventory page for a group, click its name.

Configuration Properties for a Server

The Configuration Properties section of the Inventory page for a resource contains the configuration options and currently selected values for the resource. Configuration properties for the resource are presented in three sections:

- Shared — The properties in this section relate to more than one management function, for instance both monitoring and control actions.
- Monitoring — These properties set options related to log tracking, configuration tracking, and, for servers, here is where you can disable auto-discovery of child services.
- • Control — If vRealize Hyperic supports control actions for the currently selected resource type, this section includes any configuration properties required to enable resource control, for example, the path of the start script to use to start a server.

Navigate to the Inventory Page for a Service

This topic provides instructions for accessing the Inventory page for a service

Procedure

- 1 Click **Resources > Browse > Services** to navigate to the platform.
- 2 Click the **Inventory** tab.

Viewing Inventory Properties for a Service

This topic describes the data and controls that appear at the top of the Inventory page for a service.

Resource Properties

The properties at the top of the Inventory page for a service provide identifying information including:

- Description
- Owner - By default, the vRealize Hyperic user under whose account the resource was added to inventory. Click **Change...** to assign a different resource owner.

The plugin developer controls which resource properties are displayed at the top of a resource's Inventory page: any properties enclosed in a <properties> element for a resource type appear in the page header when you browse to an instance of that type.

Map Controls for a Service

The Map control presents graphical view of the service and the resources that are related to it. The map illustrates hierarchical inventory relationships, and a resource's membership in groups or applications.

Tools Menu for a Service

When a service is selected, the **Tools** menu has the following commands:

- **Configure Service** — Opens the Configuration Properties page for the service, where you can edit the resource's configuration properties.
- **Delete Service**
- **Add to Dashboard Favorites** — Adds the service to the Favorite Resources portlet on the current Hyperic user's Dashboard.
- **Add to Group** — Opens the Group Manager page, which lists the groups to which the platform may be added. You can add a platform to a group if:
 - You have permission to access the group, and:
 - The group is a mixed group that contains platforms, servers, and services.
 - The group is a compatible group of the selected platform's platform type.

General Properties for a Service

The General Properties section of the Inventory page for a service lists the following information:

- Description
- Date Created
- Location
- Date Modified
- Resource Type

Groups Containing a Service

The Groups Containing this Resource section lists the following information for each the group of which the service is a member:

- Group
- Description

There are two controls available:

- Add to List
- Remove from List

Configuration Properties for a Service

This section displays the configuration properties for the service:

- **Shared** - These properties vary by service type.

- Monitoring - These properties control log and configuration tracking for the service.

Navigate to the Inventory Page for a Group

This topic provides instructions for accessing the Inventory page for a group.

Procedure

- 1 Click **Resources > Browse > Compatible Groups/Clusters** to navigate to the group.
- 2 Click the **Inventory** tab.

Viewing Inventory Properties for a Group

This topic describes the data and controls that appear at the top of the Inventory page for a group.

Resource Properties for a Group

The properties at the top of the Inventory page for either a compatible or mixed group are:

- Description
- Owner - By default, the vRealize Hyperic user under whose account the resource was added to inventory. Click **Change** to assign a different resource owner.

Map Control for a Compatible Group

The Map control is present on the Inventory tab for a compatible group, not for a mixed group. The map for a compatible group shows the members of the group.

Tools Menu for a Group

The **Tools** menu for a compatible or a mixed group has the following commands:

- New Group
- Delete Group
- Add to Dashboard Favorites — Adds the group to the Favorite Resources portlet on the current Hyperic user's Dashboard.
- Add to Group — Opens the Group Manager page, which lists the groups to which the platform may be added. You can add a platform to a group if:
 - You have permission to access the group, and:
 - The target group is a mixed group that contains groups.
- Schedule Downtime — This option opens the Schedule Downtime popup, where you can schedule a period of time during which alerts for members of the group will not fire.

General Properties for a Platform

The General Properties section of the Inventory page for a platform lists the following information:

- Description
- Date Created
- Location
- Date Modified
- Resource Type

Resources

This section lists the following information for each the group of which the service is a member:

- Group
- Description

There are two controls available:

- Add to List
- Remove from List

Roles a Group is Assigned To

The Roles Assigned To section lists the roles to which the group is assigned, which, along with the permission matrix for the role, governs the access that users with the role have to resources in the group.

Navigate to the Inventory Page for an Application

In vRealize Hyperic, an application is an inventory type, configured by an authorized user. An application is a set of selected services, usually running in different servers on multiple platforms, that together fulfill a single business purpose.

Configuring applications enables you to manage your infrastructure from an application - as opposed to a hardware - perspective.

Procedure

- 1 Click **Resources > Browse > Applications** to navigate to the application.
- 2 Click the **Inventory** tab.

This is the tab you use to add services to an application.

The Service Counts section shows the total number of services in the application, and the number of each type.

The Services section lists key information for each service in the application.

Viewing Inventory Properties for an Application

This topic describes the data and controls that appear at the top of the Inventory page for an application.

Resource Properties

The properties at the top of the Inventory page for an applications are:

- **Owner** - By default, the vRealize Hyperic user under whose account the resource was added to inventory. Click **Change** to assign a different resource owner.

Tools Menu for an Application

When an application is selected, the **Tools** menu has the following commands:

- **Delete Application** — remove the application from inventory.
- **Add to Dashboard Favorites** — Adds the application to the Favorite Resources portlet on the the current user's Dashboard.
- **Add to Group** — Opens the Group Manager page.

General Properties for an Application

The General Properties section of the Inventory page for an application lists the following information:

- **Description**
- **Date Created**
- **Location**
- **Date Modified**

Application Properties

The General Application section of the Inventory page for an application lists the following information about the application:

- **Application Type** - This property, supported in previous versions of Hyperic, is no longer used.
- **Engineering Contact** - An optional property.
- **Business Owner** - The date the application was created.
- **IT Operations Contact** - The date the application was last modified and the user who modified it.

Service Counts for an Application

The Service Counts section of the Inventory page for an application lists the total number of services in the application and the number of services of each service type.

Services in an Application

The Services section of the Inventory page for an application lists the following information for each of the services in the application:

- Name
- Entry Point - Whether the service is the entry point for the application.
- Service Type
- Res Type
- Host Server - Name of the server where the service runs.
- Availability

You can update the application in one of the following ways:

- To add services to the application, click **Add to List**--the Edit ApplicationName appears.
- To remove one or more services from the application, click the box to the left of each service you wish to delete, and click **Delete**.

Groups Containing an Application

The Groups Containing This Resource section of the Inventory page for an application lists the groups to which it belongs:

- To view the inventory page for a group to which the application belongs, click the group's name.
- Click **Add To List** to select the groups to which you wish to add the application.
- To remove the application from a group, checkmark the group, and click **Remove From List**.

Group Manager

The Group Manager, which appears when you click **Add to List** on the Groups Containing This Resource section of a resource's Inventory page, lists the groups to which you can add the application. Only groups that contain applications appear on the list, and only those that do not already contain the current application.

Creating Network Entities Manually

Most resource types are automatically discovered, and auto-discovery is the only way an operating system platform can be added to inventory. Conversely, auto-discovery of network or virtual platform types is not possible — you must add such platforms to inventory manually.

Most server types are auto-discovered. You create a server manually only if the managing plugin does not support server discovery, or if the plugin fails to discover a server, for instance, because it has an other-than-expected name in the process list or Windows registry.

Network services that are monitored remotely cannot be auto-discovered—you must manually create a platform service to represent a network service.

- [Create a Platform](#)

This topic provides instructions for creating a platform in vRealize Hyperic.

- [Create a Platform Service](#)

This topic provides instructions for creating a platform service.

- [Configuration Instructions for Platform Services](#)

This topic lists all supported platform service types in vRealize Hyperic.

- [Create a Server](#)

This topic provides instructions for creating a server.

- [Create a Service](#)

This topic provides instructions for creating a service.

Create a Platform

This topic provides instructions for creating a platform in vRealize Hyperic.

Operating system platform types are auto-discovered and cannot be manually added to inventory. The only platform types you can manually add to inventory are:

- Cisco IOS
- Cisco PIXOS
- GemFire Distributed System
- NetApp Filer
- Network Device
- VMWare VI3 Host
- Network Host
- VMWare vSphere Host
- VMWare vSphere VM

- Xen Host

To add a platform to the inventory using the user interface:

Procedure

- 1 Click **New Platform** on the **Tools** menu under **Browse > Resource**.
- 2 On the New Platform page, enter the following:
 - Name
 - Description
 - Location
 - Platform Type - Select the platform type from the list. Once you create the platform, you cannot change its type.
 - Full Qualified Domain Name
 - Agent Connection - The IP address:port pair of the agent the proxied connection to this platform should go through. When a platform is created manually (on this screen), the platform cannot or does not run an agent itself, and therefore the agent must connect to the platform via proxy. The list contains the IP address:port pairs from all the currently deployed Agents.
 - IP Address
 - MAC Address
 - Netmask
- 3 Click **OK**.

What to do next

Configure Platform for Monitoring - After creating the new platform, you may need supply values for one or more configuration properties. See the resource's Inventory page for configuration options and requirements.

Create a Platform Service

This topic provides instructions for creating a platform service.

There are two types of platform services in vRealize Hyperic:

- System services local to an operating system platform, such as CPUs, network interfaces, filesystems, and so on. Most local service are auto-discovered by the system plugin.
- Remote services that the vRealize Hyperic agent monitors over the network, such as HTTP, FTP, DNS, and other services using a supported protocol. Because such services are monitored by a remote agent, they cannot be auto-discovered. For network services, you manually configure platform service on the platform running the agent that will monitor the service.

To create a platform service:

Procedure

- 1 Select **Browse > Resources > Platforms** to navigate to the platform to which you wish to add a service.
- 2 Select **New Platform Service** from the **Tools** menu.
- 3 On the New Platform Service page enter:
 - Name
 - Description
 - Service Type
- 4 Click **OK**.
The new service is created.
The Inventory tab for the new service appears and prompts you to configure the resource.
- 5 Click **Configuration Properties** in the prompt.
- 6 On the Configuration Properties page, enter values for the required configuration properties, which are prefixed with a red asterisk, and optional properties as appropriate. For information about the configuration properties for a service type, refer to the Configuration Notes column for the respective property in [Configuration Instructions for Platform Services](#).

Configuration Instructions for Platform Services

This topic lists all supported platform service types in vRealize Hyperic.

The Configuration Notes column for a network service or another platform service type that require configuration has a link to a configuration properties reference for the service type. No configuration instructions are provided for auto-discovered platform services.

Table 14-1. Supported Platform Service Types

Service	Description	Configuration Notes
CPU	The vRealize Hyperic agent on a platform automatically discovers CPUs on the platform.	It is not necessary to manually configure CPU monitoring. For information about metrics available for a CPU, see <i>CPU Service in vRealize Hyperic Resource Configuration and Metrics</i> .
DHCP	Use to monitor a remote Dynamic Host Configuration Protocol server.	See <i>DHCP Platform Service in vRealize HypericResource Configuration and Metrics</i> .
DNS	Use to monitor a remote Domain Name System server.	See <i>DNS Platform Service in vRealize Hyperic Resource Configuration and Metrics</i> .
FileServer Directory	Use to monitor a directory.	See <i>FileServer Directory Service in vRealize HypericResource Configuration and Metrics</i> .

Table 14-1. Supported Platform Service Types (Continued)

Service	Description	Configuration Notes
FileServer File	Use to monitor a file.	See <i>FileServer File Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
FileServer Mount	Use to monitor a filesystem mount point and associated disks and raid arrays. Note The vRealize Hyperic Agent auto-discovers local mount points. You only need to explicitly configure a FileServer Mount service to monitor a remote network file system (NFS). Alternatively, you can install an agent on the system that hosts the NFS, in which case the NFS will be auto-discovered.	See <i>FileServer Mount Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
FileServer Directory Tree	Use to monitor a directory and the entire tree under that directory.	See <i>FileServer Directory Tree Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
HTTP	Use to monitor a specific URL.	See <i>HTTP Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
FTP	Use to monitor a specific file transfer protocol server.	See <i>FTP Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
IMAP	Use to monitor a remote Internet Message Access Protocol server.	See <i>IMAP Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
InetAddress Ping	Use to ping a remote host for availability.	See <i>InetAddress Ping Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
LDAP	Use to monitor a remote Lightweight Directory Access Server.	See <i>LDAP Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
Multiprocess	Use to monitor multiple related processes. For example, to monitor the number of httpd processes running on Apache and the system resource they consume in aggregate.	See <i>Multiprocess Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
NetworkServer Interface	Use to monitor a network interface. Note Because the vRealize Hyperic agent auto-discovers network interfaces, manual configuration of a NetworkServer Interface is rare.	See <i>NetworkServer Interface Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .
NTP	Use to monitor a network time protocol server.	See <i>NTP Platform Service</i> in <i>vRealize Hyperic Resource Configuration and Metrics</i> .

Table 14-1. Supported Platform Service Types (Continued)

Service	Description	Configuration Notes
POP3	Use to monitor a remote Post Office Protocol 3 server. Configure along with an SMTP service to monitor incoming and outgoing email services.	See <i>POP3 Platform Service in vRealize HypericResource Configuration and Metrics</i> .
Process	Use to monitor a process using a vRealize Hyperic SIGAR Process Table Query Language (PTQL) query. To configure, you supply the PTQL query in the form: <code>Class.Attribute. operator=value</code>	See <i>Process Service in vRealize HypericResource Configuration and Metrics</i> .
RPC	Use to monitor a Remote Procedure Call service. Note Not available on Windows platforms.	See <i>RPC Platform Service in vRealize HypericResource Configuration and Metrics</i> .
Script	Used to configure vRealize Hyperic to periodically run a script that collects a system or application metric.	Script Service
SMTP	Use to monitor a remote Simple Mail Transfer Protocol server. Configure along with a POP3 service to monitor incoming and outgoing email services.	See <i>SMTP Platform Service in vRealize Hyperic Resource Configuration and Metrics</i> .
SNMP	Use to monitor a remote Simple Network Management Protocol agent.	See <i>SNMP Platform Service in vRealize Hyperic Resource Configuration and Metrics</i> .
SSH	Use to monitor a remote SSH service.	See <i>SSH Platform Service in vRealize Hyperic Resource Configuration and Metrics</i> .
TCP Socket	Use to monitor the availability of a remote TCP socket.	See <i>TCP Socket Platform Service in vRealize Hyperic Resource Configuration and Metrics</i> .
Windows Service	Use to monitor an application that runs as a service under Windows.	To configure it, you supply its Service Name in Windows. To determine the Service Name: <ol style="list-style-type: none">1 Select Start > Run from the Windows menu.2 Type <code>services.msc</code> in the run dialog and click OK.3 In the list of services displayed, right-click the service you wish to monitor and click Properties.4 Locate the Service Name on the General tab. See Windows Service Service.

Create a Server

This topic provides instructions for creating a server.

Most server types that vRealize Hyperic manages are auto-discovered, rather than manually added to inventory on the New Server page.

You may need to create a server manually under some circumstances, for instance, if:

- The plugin that manages the server does not support auto-discovery of server instances.
- The auto-discovery method used by the plugin failed, because the entry it looked for in the process table or Windows registry was not found. For example, the server's name in the process list is different than the plugin uses to detect server instances.

To create a server:

Procedure

- 1 Select **Browse > Resources > Platforms** to navigate to the platform to which you wish to add a server.
- 2 Select **New Server** from the **Tools** menu.
- 3 On the New Server page enter:
 - Name
 - Description
 - Server Type
 - Install Path - If there are other servers of the type you are creating on the current platform, make sure that the installation path you define for the new server instance is unique, and not the same as any of the other servers of the same type. For autodiscovered servers, the plugin sets the value of the `Install Path` property, and uses it as the basis for the resource's autoinventory identifier.

Note If you create a server manually, you must specify the `Install Path` property, but it need not be the actual installation path for the server. The only requirement is that all servers of the same resource type on the same platform have unique `Install Path` property values.

- 4 Click **OK**.

Inventory Page for a New Server

This topic provides information about the Inventory page for a server that has just been added to the inventory.

The inventory page contains the following sections:

- **General Properties** -The server's owner is the username for the vRealize Hyperic account that created the server.

- **Type and Host Properties** -The Host Platform is the name of the platform to which you added the server.
- **Services** -This section of the page lists services discovered in the server instance. For an new server instance, no services appear until any required configuration properties are defined. You can manually add services to a server, but it is uncommon.
- **Groups Containing this Resource** - You can click Add to List in this section to assign the server to one or more resource groups.
- **Configuration Properties** -The properties in the section vary, depending on the resource type of the server you are created.
 - **Shared** - For some server types, you need to supply resource connection or authorization credential data in this section in order to monitor and manage the resource.
 - **Monitoring** - If the server type supports log or configuration tracking, you can configure tracking options in this section.
 - **Control** - If the server type supports control actions, you configure associated options in this section.

After creating the new server, you may need supply values for one or more configuration properties. See the resource's Inventory page for configuration options and requirements.

Create a Service

This topic provides instructions for creating a service.

In Hyperic, a service is a component that runs in a server (for example, an EJB), or is associated with an operating system (for example, a CPU). System-level service, and services that an agent will manage remotely are referred to as platform services.

Most service types are auto-discovered by the plugin that manages that service type. Services that run on a platform are discovered along with the platform itself by the Hyperic's system plugin. Services that run in a local server are auto-discovered by the plugin that manages the host server — for example, the plugin that manages Tomcat discovers the services running in Tomcat.

The only service types in Hyperic that are routinely created explicitly are remote platform services.

To add a service to inventory using the user interface:

Procedure

- 1 Select **Browse > Resources > Server** to navigate to the server to which you wish to add a service.
- 2 Select **New Service** from the **Tools** menu.
- 3 On the New Service page enter:
 - Name
 - Description
 - Service Type

- Host Server

4 Click **OK**.

Remove a Resource from the Inventory

This topic provides instructions for removing a resource from the vRealize Hyperic inventory.

To remove a resource:

Procedure

- 1 Select **Browse > Resources** to navigate to the resource you wish to delete.
- 2 Select the **Delete ResourceType** option from the **Tools** menu.

Monitoring ceases and the resource is removed from inventory. If the resource has child resources, those children (and their children, as applicable) will be removed as well.

Resources you delete will disappear from the resource list immediately. Removal of the resources from the database occurs asynchronously; there may be a slight delay before resources are removed from the database. If you restart the agent prior to the resource being removed from the database, errors might occur if you re-add the resource to memory after the agent rediscovers it.

Do not try to re-import a platform immediately after deleting it — this may fail. Also, until the delete process running in background completes, a deleted platform will still appear in the vRealize Hyperic Health's agent tab. Similarly, the vRealize Hyperic license count displayed in the web user interface will not be decremented until the delete process is completed.

Deploying and Managing Plugins

15

vRealize Hyperic product plugin management and deployment features ensure that all agents reporting to the server run exactly the same product plugins as those deployed to the server.

The key features are:

- Plugin management interface - Plugin Manager is a user interface for deployment-wide management of product plugins. You use the Plugin Manager to deploy and manage plugins on the server and agents that report to the server.
- Server-Agent Plugin Synchronization (SAPS) - An agent, and the server to which it reports, cooperate to ensure that the agent is running exactly the same set of plugins as the server.

Version Requirements for Plugin Synchronization

The SAPS process will not synchronize plugins on agents of an earlier version than the server. The server can only synchronize plugins on agents running the same or later version. The Plugin Manager user interface indicates how many of the agents reporting to the server are in sync. If there are agents that are out-of-date, you can see which agents need upgrading.

This section includes the following topics:

- [Plugin Administration Tasks](#)
- [Resource Plugin Deployment Directories](#)
- [Working with Plugin Manager](#)
- [About Custom Plugin Deployment with Plugin Manager](#)
- [Removing a Plugin with Plugin Manager](#)
- [Synchronizing Plugins at Startup](#)
- [Disable Server-Agent Plugin Synchronization](#)

Plugin Administration Tasks

If you run vRealize Hyperic with SAPS enabled, you can use Plugin Manager to administer plugins on up-to-date agents throughout your deployment.

The primary tasks you perform are:

- Deploy and update custom plugins — If you develop your own plugin, or customize a built-in plugin, use Plugin Manager to upload and deploy the new or modified plugin to the server and to all up-to-date agents in your deployment. Plugins you deploy with Plugin Manager are deployed to a directory external to the server's installation directory, so that they are not overwritten by a server upgrade.
- Remove plugins — For optimal vRealize Hyperic scalability, it is good practice to reduce the agent footprint by running only the plugins you use, and to remove a plugin if you have no resources of the types the plugin manages. You can remove one or more plugins from all up-to-date agents in your deployment with a single Plugin Manager command.

Since version 4.6, vRealize Hyperic you cannot manually deploy your plugins.

Resource Plugin Deployment Directories

This topic provides information about the resource plugin deployment directories.

There are two product plugin deployment directories on the vRealize Hyperic server:

- Standard plugins - the *ServerHome/hq-engine/hq-server/webapps/ROOT/WEB_INF/hq-plugins/* directory is used for standard plugins packaged with vRealize Hyperic server.
- Custom plugins - the *user.dir/hq-plugins* directory is used for custom plugins or a customized version of a built-in plugin. Plugin Manager deploys the plugin to the *hq-plugins* subdirectory of the current working directory of the server process. Custom plugins are deployed to a separate plugin directory, external to the server installation, so that they will not be overwritten when you upgrade the vRealize Hyperic server.

You can configure the server to store custom plugins in a different directory with the `server.plugin.custom.dir` property.

On the Hyperic Agent all product plugins reside in the *AgentHome/bundles/agent-x.y.z/pdk/plugins/* directory.

Working with Plugin Manager

Plugin Manager, available on the Server Settings section of the Administration tab, is a user interface for deployment-wide management of product plugins.

You can use the Plugin Manager to administer plugins on the vRealize Hyperic server and on all up-to-date agents that report to the server.

Plugin Manager displays all plugins deployed to the server, and the status of each plugin on agents reporting to the server.

You can use Plugin Manager to upload new plugins, update existing plugins, and remove plugins. The changes you make are replicated on all up-to-date agents that report to the server. Note that plugins you upload must comply with plugin naming conventions, and may not be larger than 5 MB.

When you remove one or more plugins using Plugin Manager, all resources in inventory of types managed by the plugin are removed from inventory.

About Custom Plugin Deployment with Plugin Manager

This topic provides information about what happens when you upload a plugin with Plugin Manager.

When you upload a new or updated plugin with the Plugin Manager:

- The plugin is copied to the vRealize Hyperic server's `user.dir/hq-plugins` directory.
- If a product plugin of the same name exists in the server's `hq-plugins` directory (likely to be the case if you are deploying a customized version of a built-in plugin), the version in `user.dir/hq-plugins` henceforth is loaded, rather than the version in the server's `hq-plugins` directory (`ServerHome/hq-engine/hq-server/webapps/ROOT/WEB_INF/hq-plugins/`).
- The vRealize Hyperic server updates and restarts each up-to-date agent. Note that agent update and agent restart are asynchronous - the server processes two queues: one of agents that still need the update pushed, and one of those that have the update and need to be restarted.
 - The server pushes the new or changed plugins to all up-to-date agents, up to four at a time. A built-in throttle mechanism limits the number of agent update threads.
 - The server issues a restart command to each updated agent, up to 20 in parallel. A built-in throttle mechanism limits the number of agent restart threads.

Removing a Plugin with Plugin Manager

You can use the Plugin Manager to remove one or more plugins from the vRealize Hyperic server and all up-to-date agents that report to the server.

When you remove a plugin:

- The plugin is removed from its deployment directory on the vRealize Hyperic server.
- The plugin is removed from its deployment directory on each up-to-date vRealize Hyperic agent reporting to the server, and each agent is restarted.
- All resources in inventory managed by the removed plugin are removed from vRealize Hyperic's resource inventory.

Synchronizing Plugins at Startup

Server-side and agent-side plugins can become out-of-sync if you manually add or remove plugins and do not apply updates consistently to the server and to every agent.

To detect plugin mismatches caused by inconsistent manual deployment, the vRealize Hyperic server checks for out-of-sync plugins every time it starts up, and every time an agent connects to it upon agent start or restart. (The server does the post-startup checks five minutes after startup, so that any agents that were restarted while the server was down to sync their plugin inventory.) In the case of a mismatch between server-side and agent-side plugins, the server updates the agent.

Synchronizing Plugins at Agent Startup

When an up-to-date agent starts up and loads the plugins in its plugin directory, it sends a plugin status report to the server, including the MD5 checksum of each plugin it loaded.

On receipt of a plugin status report from an up-to-date agent, the vRealize Hyperic server compares the MD5 checksum of each plugin that the agent loaded with the MD5 checksum of the corresponding plugin (a plugin with the same name) currently deployed on the server.

For a mismatch:

- If the agent loaded a plugin with the same name but different MD5 checksum than one that is currently deployed to the server, the server pushes its version to the agent.
- If the agent did not load a plugin that is currently deployed to the server (the agent did not load a plugin with the same name as one that the server loaded), the server pushes the plugin to the agent.
- If the agent loaded a plugin that is not deployed to the vRealize Hyperic server (there is no plugin with the same name on the server), the server removes the plugin from the agent.

After adding or removing agent plugins as necessary, the server restarts the agent.

Synchronizing Plugins at Server Startup

When the vRealize Hyperic server starts up and loads the plugins in its deployment directories, it compares the MD5 checksum of the plugins it loaded with the MD5 checksums in the plugin status report last received from each agent.

For a mismatch:

- If an agent's last plugin status report included a plugin with the same name but different MD5 checksum than one that is currently deployed to the vRealize Hyperic server, the server pushes its version to the agent.
- If an agent's last plugin status report does not include a plugin that is currently deployed to the server (the agent did not load a plugin with the same name as one that the server loaded), the server pushes the plugin to the agent.

- If an agent's last plugin status report included a plugin that is not deployed to the server (there is no plugin with the same name on the server), the server removes the plugin from the agent.

After updating the agent plugins appropriately, the server restarts the agent.

Disable Server-Agent Plugin Synchronization

This topic provides instructions for disabling server-agent plugin synchronization.

SAPS is enabled by default. To disable SAPS:

Procedure

- ◆ Add the following line to the *ServerHome/conf/hq-server.conf* file.

```
server.pluginsync.enabled=false
```