

vRealize Operations Manager Configuration Guide

06 NOV 2020

vRealize Operations Manager 7.0

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

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

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

Copyright © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Configuration 8

1 Connecting to Data Sources 9

VMware vSphere Solution 9

Configure a vCenter Adapter Instance 11

Configure User Access for Actions 12

End Point Operations Management Solution 13

End Point Operations Management Agent Installation and Deployment 14

Roles and Privileges 57

Registering Agents on Clusters 58

Manually Create Operating System Objects 59

Managing Objects with Missing Configuration Parameters 60

Mapping Virtual Machines to Operating Systems 61

Customizing How End Point Operations Management Monitors Operating Systems 62

VMware vRealize Application Management Pack 73

View the Configuration Details 74

Log Insight 74

Log Insight Page 74

Logs Tab 75

Configuring vRealize Log Insight with vRealize Operations Manager 75

Log Forwarding 77

Business Management 78

Configure the vRealize Business for Cloud Adapter 79

Cost Settings for Financial Accounting Model 79

Overview of Cost Drivers 81

Editing Cost Drivers 83

Cluster Cost Overview 89

Cost Calculation Status Overview 89

vRealize Automation Solution 90

Supported vRealize Automation Versions 90

Object Types and Relationships 91

vRealize Automation Workload Placement 91

Port Information 92

Security Guidelines 92

Configuring vRealize Automation 93

Alert Definitions 96

vSAN 97

Configure a vSAN Adapter Instance 97

Verify that the Adapter Instance is Connected and Collecting Data	99
Installing Optional Solutions	100
Managing Solution Credentials	101
Managing Collector Groups	102
2 Configuring Alerts and Actions	103
Types of Alerts	103
Configuring Alerts	103
Defining Alerts in vRealize Operations Manager	103
Defining Symptoms for Alerts	104
Defining Recommendations for Alert Definitions	108
Create a New Alert Definition	109
Alert Definition Best Practices	110
Creating and Managing Alert Notifications	111
Create an Alert Definition for Department Objects	124
Alerts Group	136
Viewing Actions	137
List of vRealize Operations Manager Actions	137
Actions Supported for Automation	139
Integration of Actions with vRealize Automation	140
Working with Actions That Use Power Off Allowed	141
3 Configuring and Using Workload Optimization	145
Configuring Workload Optimization	146
Business Intent: Tag-Based VM Placement in Clusters	147
Business Intent - Host-Based Virtual Machine Placement	150
Business Intent Workspace	151
Configuring Workload Optimization Alerts	152
Using Workload Optimization	153
Example: Run Workload Optimization	154
Example: Schedule a Repeating Optimization Action	155
Example: Run Workload Optimization from Recommended Actions	157
4 Configuring Policies	159
Policies	159
Policy Decisions and Objectives	161
Active Policies Tab for Policies	161
Policy Library Tab for Policies	164
Operational Policies	166
Types of Policies	167
Custom Policies	167

Default Policy in vRealize Operations Manager	169
Policies Provided with vRealize Operations Manager	169
Using the Monitoring Policy Workspace to Create and Modify Operational Policies	171
Policy Workspace in vRealize Operations Manager	172
5 Configuring Super Metrics	189
Create a Super Metric	190
Enhancing Your Super Metrics	192
Exporting and Importing a Super Metric	193
6 Configuring Objects	195
Object Discovery	195
About Objects	196
Managing Objects in Your Environment	198
Managing Custom Object Groups	204
Managing Application Groups	208
7 Configuring Data Display	211
Widgets	211
Widget Interactions	212
Manage Metric Configuration	212
Add a Resource Interaction XML File	213
Widget Definitions List	215
Dashboards	217
Types Of Dashboards	217
Create and Configure Dashboards	238
Managing Dashboards	241
Views	246
Views Overview	247
Views and Reports Ownership	248
Create and Configure a View	248
Editing, Cloning, and Deleting a View	261
User Scenario: Create, Run, Export, and Import a vRealize Operations Manager View for Tracking Virtual Machines	262
Reports	264
Report Templates Tab	265
Generated Reports Tab	265
Create and Modify a Report Template	266
Add a Network Share Plug-In for vRealize Operations Manager Reports	269
8 Configuring Application Monitoring with Wavefront	271
Configure a Wavefront Account	272

Configure the Wavefront Account	272
Create a Wavefront Trial Account	273
Manage VMware Application Proxy	273
Deploy VMware Application Proxy	275
Configure Network Time Protocol Settings	277
Add and Configure an Application Proxy	278
Manage Agents in Virtual Machines	279
Install an Agent	282
Manage Application Services	283
Activate and Deactivate an Application Service	284
Uninstall an Agent	285
Monitor Metrics in Wavefront	286

9 Configuring Administration Settings 287

Managing Users and Access Control	287
Users of vRealize Operations Manager	288
Roles and Privileges	292
User Scenario: Manage User Access Control	293
Configure a Single Sign-On Source	296
Audit Users and the Environment	299
Passwords and Certificates	301
Reset the Administrator Password	301
Generate a Passphrase	302
Custom Certificates	302
Modifying Global Settings	308
List of Global Settings	308
Global Settings	311
Transfer Ownership of Dashboards and Report Schedules	312
Create a Support Bundle	313
Customizing Icons	313
Customize an Object Type Icon	314
Customize an Adapter Type Icon	314

10 OPS-CLI Command-Line Tool 316

dashboard Command Operations	317
template Command Operations	318
supermetric Command Operations	319
attribute Command Operations	320
reskind Command Operations for Object Types	320
report Command Operations	320
view Command Operations	321

[file Command Operations](#) 321

About Configuration

The VMware *vRealize Operations Manager Configuration Guide* describes how to configure and monitor your environment. It shows you how to connect vRealize Operations Manager to external data sources and analyze the data collected from them, ensure that users and their supporting infrastructure are in place, configure resources to determine the behavior of your objects, and format the content that appears in vRealize Operations Manager.

To help you maintain and expand your vRealize Operations Manager installation, this information describes how to manage nodes and clusters, configure NTP, view log files, create support bundles, and add a maintenance schedule. It provides information about license keys and groups, and shows you how to generate a passphrase, review the certificates used for authentication, run the describe process, and perform advanced maintenance functions.

Intended Audience

This information is intended for vRealize Operations Manager administrators, virtual infrastructure administrators, and operations engineers who install, configure, monitor, manage, and maintain the objects in your environment.

For users who want to configure vRealize Operations Manager programmatically, the VMware vRealize Operations Manager REST API documentation is available in HTML format and is installed with your vRealize Operations Manager instance. For example, if the URL of your instance is `https://vrealize.example.com`, the API reference is available from `https://vrealize.example.com/suite-api/docs/rest/index.html`.

Connecting vRealize Operations Manager to Data Sources

1

Configure solutions in vRealize Operations Manager to connect to and analyze data from external data sources in your environment. Once connected, you use vRealize Operations Manager to monitor and manage objects in your environment.

A solution might be only a connection to a data source, or it might include predefined dashboards, widgets, alerts, and views.

vRealize Operations Manager includes the VMware vSphere and End Point Operations Management solutions. These solutions are installed when you install vRealize Operations Manager.

Other solutions can be added to vRealize Operations Manager as management packs, such as the VMware Management Pack for NSX for vSphere. To download VMware management packs and other third-party solutions, visit the VMware Solution Exchange at <https://marketplace.vmware.com/vsx/>.

This chapter includes the following topics:

- [VMware vSphere Solution in vRealize Operations Manager](#)
- [End Point Operations Management Solution in vRealize Operations Manager](#)
- [VMware vRealize Application Management Pack](#)
- [Log Insight](#)
- [Business Management](#)
- [vRealize Automation Solution](#)
- [vSAN](#)
- [Installing Optional Solutions in vRealize Operations Manager](#)

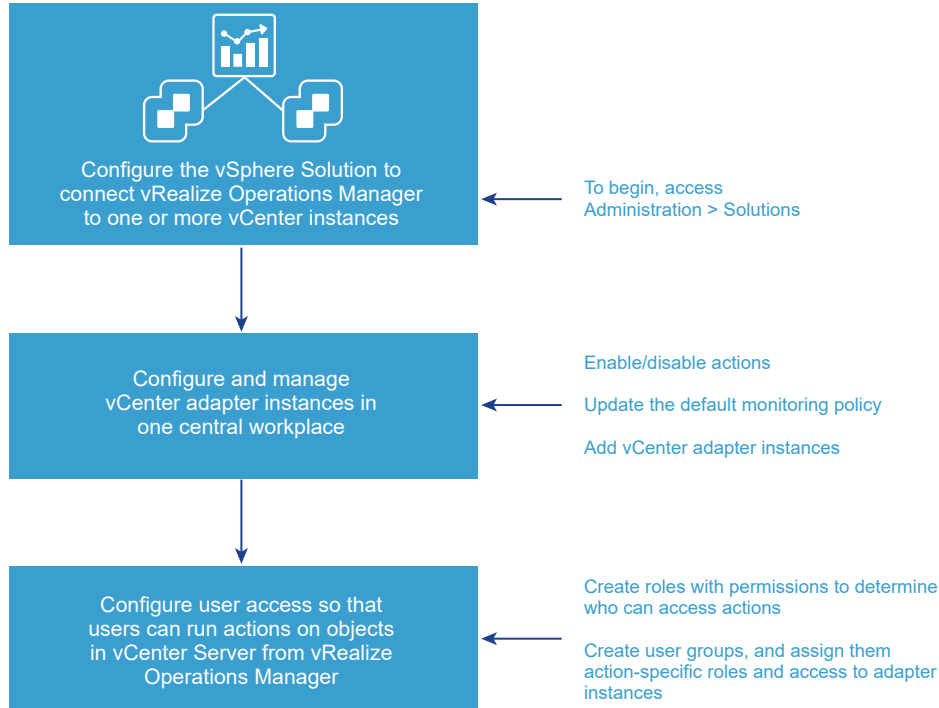
VMware vSphere Solution in vRealize Operations Manager

The VMware vSphere solution connects vRealize Operations Manager to one or more vCenter Server instances. You collect data and metrics from those instances, monitor them, and run actions in them.

vRealize Operations Manager evaluates the data in your environment, identifying trends in object behavior, calculating possible problems and future capacity for objects in your system based on those trends, and alerting you when an object exhibits defined symptoms.

Configuring the vSphere Solution

The vSphere solution is installed together with vRealize Operations Manager. The solution provides the vCenter Server adapter which you must configure to connect vRealize Operations Manager to your vCenter Server instances.



How Adapter Credentials Work

The vCenter Server credentials that you use to connect vRealize Operations Manager to a vCenter Server instance, determines what objects vRealize Operations Manager monitors. Understand how these adapter credentials and user privileges interact to ensure that you configure adapters and users correctly, and to avoid some of the following issues.

- If you configure the adapter to connect to a vCenter Server instance with credentials that have permission to access only one of your three hosts, every user who logs in to vRealize Operations Manager sees only the one host, even when an individual user has privileges on all three of the hosts in the vCenter Server.
- If the provided credentials have limited access to objects in the vCenter Server, even vRealize Operations Manager administrative users can run actions only on the objects for which the vCenter Server credentials have permission.
- If the provided credentials have access to all the objects in the vCenter Server, any vRealize Operations Manager user who runs actions is using this account.

Controlling User Access to Actions

Use the vCenter Server adapter to run actions on the vCenter Server from vRealize Operations Manager. If you choose to run actions, you must control user access to the objects in your vCenter Server environment. You control user access for local users based on how you configure user privileges in vRealize Operations Manager. If users log in using their vCenter Server account, then the way their account is configured in vCenter Server determines their privileges.

For example, you might have a vCenter Server user with a read-only role in vCenter Server. If you give this user the vRealize Operations Manager Power User role in vCenter Server rather than a more restrictive role, the user can run actions on objects because the adapter is configured with credentials that has privileges to change objects. To avoid this type of unexpected result, configure local vRealize Operations Manager users and vCenter Server users with the privileges you want them to have in your environment.

Configure a vCenter Adapter Instance in vRealize Operations Manager

To manage your vCenter Server instances in vRealize Operations Manager, you must configure an adapter instance for each vCenter Server instance. The adapter requires the credentials that are used for communication with the target vCenter Server.

Caution Any adapter credentials you add are shared with other adapter administrators and vRealize Operations Manager collector hosts. Other administrators might use these credentials to configure a new adapter instance or to move an adapter instance to a new host.

Prerequisites

Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data. If the provided credentials have limited access to objects in vCenter Server, all users, regardless of their vCenter Server privileges see only the objects that the provided credentials can access. At a minimum, the user account must have Read privileges and the Read privileges must be assigned at the data center or vCenter Server level.

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Solutions**.
- 2 On the Solutions page, select **VMware vSphere** and click the **Configure** icon.
- 3 Enter a display name and description for the adapter instance.
- 4 In the **vCenter Server** text box, enter the FQDN or IP address of the vCenter Server instance to which you are connecting.

The vCenter Server FQDN or IP address must be reachable from all nodes in the vRealize Operations Manager cluster.

- 5 To add credentials for the vCenter Server instance, click the **Add** icon, and enter the required credentials. The vCenter credential must have **Performance > Modify intervals** permission enabled in the target vCenter to collect VM guest metrics.

- 6 The adapter is configured to run actions on objects in the vCenter Server from vRealize Operations Manager. If you do not want to run actions, select **Disable**.

The credentials provided for the vCenter Server instance are also used to run actions. If you do not want to use these credentials, you can provide alternative credentials by expanding **Alternate Action Credentials**, and clicking the **Add** icon.

- 7 Click **Test Connection** to validate the connection with your vCenter Server instance.

- 8 In the **Review and Accept Certificate** dialog box, review the certificate information.

- ◆ If the certificate presented in the dialog box matches the certificate for your target vCenter Server, click **OK**.
- ◆ If you do not recognize the certificate as valid, click **Cancel**. The test fails and the connection to vCenter Server is not completed. You must provide a valid vCenter Server URL or verify the certificate on the vCenter Server is valid before completing the adapter configuration.

- 9 To modify the advanced options regarding collectors, object discovery, or change events, expand the **Advanced Settings**.

For information about these advanced settings, search for the VMware vSphere Solution Workspace Options in the Information Center.

- 10 To adjust the default monitoring policy that vRealize Operations Manager uses to analyze and display information about the objects in your environment, click **Define Monitoring Goals**.

For information about monitoring goals, search for the VMware vSphere Solution Workspace Options in the Information Center.

- 11 Click **Save Settings**.

The adapter instance is added to the list.

Results

vRealize Operations Manager begins collecting data from the vCenter Server instance. Depending on the number of managed objects, the initial collection can take more than one collection cycle. A standard collection cycle begins every five minutes.

What to do next

If you configured the adapter to run actions, configure user access for the actions by creating action roles and user groups.

Configure User Access for Actions

To ensure that users can run actions in vRealize Operations Manager, you must configure user access to the actions.

You use role permissions to control who can run actions. You can create multiple roles. Each role can give users permissions to run different subsets of actions. Users who hold the Administrator role or the default super user role already have the required permissions to run actions.

You can create user groups to add action-specific roles to a group rather than configuring individual user privileges.

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Access > Access Control**.
- 2 To create a role:
 - a Click the **Roles** tab.
 - b Click the **Add** icon, and enter a name and description for the role.
- 3 To apply permissions to the role, select the role, and in the Permissions pane, click the **Edit** icon.
 - a Expand **Environment**, and then expand **Action**.
 - b Select one or more of the actions, and click **Update**.
- 4 To create a user group:
 - a Click the **User Groups** tab, and click the **Add** icon.
 - b Enter a name for the group and a description, and click **Next**.
 - c Assign users to the group, and click the **Objects** tab.
 - d Select a role that has been created with permissions to run actions, and select the **Assign this role to the user** check box.
 - e Configure the object privileges by selecting each adapter instance to which the group needs access to run actions.
 - f Click **Finish**.

What to do next

Test the users that you assigned to the group. Log out, and log back in as one of the users. Verify that this user can run the expected actions on the selected adapter.

End Point Operations Management Solution in vRealize Operations Manager

You configure End Point Operations Management to gather operating system metrics and to monitor availability of remote platforms and applications. This solution is installed with vRealize Operations Manager.

End Point Operations Management Agent Installation and Deployment

Use the information in these links to help you to install and deploy End Point Operations Management agents in your environment.

Prepare to Install the End Point Operations Management Agent

Before you can install the End Point Operations Management agent, you must perform preparatory tasks.

Prerequisites

- To configure the agent to use a keystore that you manage yourself for SSL communication, set up a JKS-format keystore for the agent on its host and import its SSL certificate. Make a note of the full path to the keystore, and its password. You must specify this data in the agent's `agent.properties` file.

Verify that the agent keystore password and the private key password are identical.

- Define the agent `HQ_JAVA_HOME` location.

vRealize Operations Manager platform-specific installers include JRE 1.8.x . Depending on your environment and the installer you use, you may need to define the location of the JRE to ensure that the agent can find the JRE to use. See [Configuring JRE Locations for End Point Operations Management Components](#).

Supported Operating Systems for the End Point Operations Management Agent

These tables describe the supported operating systems for End Point Operations Management agent deployments.

These configurations are supported for the agent in both development and production environments.

Table 1-1. Supported Operating Systems for the End Point Operations Management Agent

Operating System	Processor Architecture	JVM
RedHat Enterprise Linux (RHEL) 5.x, 6.x, 7.x	x86_64, x86_32	Oracle Java SE8
CentOS 5.x, 6.x, 7.x	x86_64, x86_32	Oracle Java SE8
SUSE Enterprise Linux (SLES) 11.x, 12.x	x86_64	Oracle Java SE8
Windows 2008 Server, 2008 Server R2	x86_64, x86_32	Oracle Java SE8
Windows 2012 Server, 2012 Server R2	x86_64	Oracle Java SE8
Windows Server 2016	x86_64	Oracle Java SE8
Solaris 10, 11	x86_64, SPARC	Oracle Java SE7
AIX 6.1, 7.1	Power PC	IBM Java SE7

Table 1-1. Supported Operating Systems for the End Point Operations Management Agent (continued)

Operating System	Processor Architecture	JVM
VMware Photon Linux 1.0	x86_64	Open JDK 1.8.0_72-BLFS
Oracle Linux versions 5, 6, 7	x86_64, x86_32	Open JDK Runtime Environment 1.7

Selecting an Agent Installer Package

The End Point Operations Management agent installation files are included in the vRealize Operations Manager installation package.

You can install the End Point Operations Management agent from a `tar.gz` or `.zip` archive, or from an operating system-specific installer for Windows or for Linux-like systems that support RPM.

When you install a non-JRE version of End Point Operations Management agent, to avoid being exposed to security risks related to earlier versions of Java, it is recommended that you only use the latest Java version.

- [Install the Agent on a Linux Platform from an RPM Package](#)

You can install the End Point Operations Management agent from a RedHat Package Manager (RPM) package. The agent in the `noarch` package does not include a JRE.

- [Install the Agent on a Linux Platform from an Archive](#)

You can install an End Point Operations Management agent on a Linux platform from a `tar.gz` archive.

- [Install the Agent on a Windows Platform from an Archive](#)

You can install an End Point Operations Management agent on a Windows platform from a `.zip` file.

- [Install the Agent on a Windows Platform Using the Windows Installer](#)

You can install the End Point Operations Management agent on a Windows platform using a Windows installer.

- [Installing an End Point Operations Management Agent Silently on a Windows Machine](#)

You can install an End Point Operations Management agent on a Windows machine using silent or very silent installation.

- [Install the Agent on an AIX Platform](#)

You can install the End Point Operations Management agent on an AIX platform.

- [Install the Agent on a Solaris Platform](#)

You can install the End Point Operations Management agent on a Solaris platform.

Install the Agent on a Linux Platform from an RPM Package

You can install the End Point Operations Management agent from a RedHat Package Manager (RPM) package. The agent in the `noarch` package does not include a JRE.

Agent-only archives are useful when you deploy agents to a large number of platforms with various operating systems and architectures. Agent archives are available for Windows and UNIX-like environments, with and without built-in JREs.

The RPM performs the following actions:

- Creates a user and group named `epops` if they do not exist. The user is a service account that is locked and you cannot log into it.
- Installs the agent files into `/opt/vmware/epops-agent`.
- Installs an init script to `/etc/init.d/epops-agent`.
- Adds the init script to `chkconfig` and sets it to on for run levels 2, 3, 4, and 5.

If you have multiple agents to install, see [Install Multiple End Point Operations Management Agents Simultaneously](#).

Prerequisites

- Verify that you have sufficient privileges to deploy an End Point Operations Management agent. You must have vRealize Operations Manager user credentials that include a role that allows you to install End Point Operations Management agents. See [Roles and Privileges in vRealize Operations Manager](#).
- If you plan to run ICMP checks, you must install the End Point Operations Management agent with **root** privileges.
- To configure the agent to use a keystore that you manage yourself for SSL communication, set up a JKS-format keystore for the agent on its host and configure the agent to use its SSL certificate. Note the full path to the keystore, and its password. You must specify this data in the agent `agent.properties` file.

Verify that the agent keystore password and the private key password are identical.

- If you are installing a non-JRE package, define the agent `HQ_JAVA_HOME` location.

End Point Operations Management platform-specific installers include JRE 1.8.x. Platform-independent installers do not. Depending on your environment and the installer you use, you might need to define the location of the JRE to ensure that the agent can find the JRE to use. See [Configuring JRE Locations for End Point Operations Management Components](#).

- If you are installing a non-JRE package, verify that you are using the latest Java version. You might be exposed to security risks with earlier versions of Java.
- Verify that the installation directory for the End Point Operations Management agent does not contain a vRealize Hyperic agent installation.
- If you are using the `noarch` installation, verify that a JDK or JRE is installed on the platform.
- Verify that you use only ASCII characters when specifying the agent installation path. If you want to use non-ASCII characters, you must set the encoding of the Linux machine and SSH client application to UTF-8.

Procedure

- 1 Download the appropriate RPM bundle to the target machine.

Operating System	RPM Bundle to Download
64bit Operating System	<code>epops-agent-x86-64-linux-version.rpm</code>
32bit Operating System	<code>epops-agent-x86-linux-version.rpm</code>
No Arch	<code>epops-agent-noarch-linux-version.rpm</code>

- 2 Open an SSH connection using root credentials.
- 3 Run `rpm -i epops-agent-Arch-linux-version.rpm` to install the agent on the platform that the agent will monitor, where *Arch* is the name of the archive and *version* is the version number.

Results

The End Point Operations Management agent is installed, and the service is configured to start at boot.

What to do next

Before you start the service, verify that the `epops` user credentials include any permissions that are required to enable your plug-ins to discover and monitor their applications, then perform one of the following processes.

- Run `service epops-agent start` to start the `epops-agent` service.
- If you installed the End Point Operations Management agent on a machine running SuSE 12.x, start the End Point Operations Management agent by running the `[EP Ops Home]/bin/ep-agent.sh start` command.
- When you attempt to start an End Point Operations Management agent you might receive a message that the agent is already running. Run `./bin/ep-agent.sh stop` before starting the agent.
- Configure the agent in the `agent.properties` file, then start the service. See [Activate End Point Operations Management Agent to vRealize Operations Manager Server Setup Properties](#).

Install the Agent on a Linux Platform from an Archive

You can install an End Point Operations Management agent on a Linux platform from a `tar.gz` archive.

By default, during installation, the setup process prompts you to provide configuration values. You can automate this process by specifying the values in the agent properties file. If the installer detects values in the properties file, it applies those values. Subsequent deployments also use the values specified in the agent properties file.

Prerequisites

- Verify that you have sufficient privileges to deploy an End Point Operations Management agent. You must have vRealize Operations Manager user credentials that include a role that allows you to install End Point Operations Management agents. See [Roles and Privileges in vRealize Operations Manager](#).
- If you plan to run ICMP checks, you must install the End Point Operations Management agent with **root** privileges.
- Verify that the installation directory for the End Point Operations Management agent does not contain a vRealize Hyperic agent installation.
- Verify that you use only ASCII characters when specifying the agent installation path. If you want to use non-ASCII characters, you must set the encoding of the Linux machine and SSH client application to UTF-8.

Procedure

- 1 Download and extract the End Point Operations Management agent installation `tar.gz` file that is appropriate for your Linux operating system.

Operating System	tar .gz Bundle to Download
64bit Operating System	<code>epops-agent-x86-64-linux-version.tar.gz</code>
32bit Operating System	<code>epops-agent-x86-linux-version.tar.gz</code>
No Arch	<code>epops-agent-noJRE-version.tar.gz</code>

- 2 Run `cd agent_name/bin` to open the `bin` directory for the agent.
- 3 Run `ep-agent.sh start`.

The first time that you install the agent, the command launches the setup process, unless you already specified all the required configuration values in the agent properties file.

- 4 (Optional) Run `ep-agent.sh status` to view the current status of the agent, including the IP address and port.

What to do next

Register the client certificate for the agent. See [Regenerate an Agent Client Certificate](#).

Install the Agent on a Windows Platform from an Archive

You can install an End Point Operations Management agent on a Windows platform from a `.zip` file.

By default, during installation, the setup process prompts you to provide configuration values. You can automate this process by specifying the values in the agent properties file. If the installer detects values in the properties file, it applies those values. Subsequent deployments also use the values specified in the agent properties file.

Prerequisites

- Verify that you have sufficient privileges to deploy a End Point Operations Management agent. You must have vRealize Operations Manager user credentials that include a role that allows you to install End Point Operations Management agents. See [Roles and Privileges in vRealize Operations Manager](#).
- Verify that the installation directory for the End Point Operations Management agent does not contain a vRealize Hyperic agent installation.
- Verify that you do not have any End Point Operations Management or vRealize Hyperic agent installed on your environment before running the agent Windows installer.

Procedure

- 1 Download and extract the End Point Operations Management agent installation .zip file that is appropriate for your Windows operating system.

Operating System	ZIP Bundle to Download
64bit Operating System	epops-agent-x86-64-win-version.zip
32bit Operating System	epops-agent-win32-version.zip
No Arch	epops-agent-noJRE-version.zip

- 2 Run `cd agent_name\bin` to open the bin directory for the agent.
- 3 Run `ep-agent.bat install`.
- 4 Run `ep-agent.bat start`.

The first time that you install the agent, the command starts the setup process, unless you already specified the configuration values in the agent properties file.

What to do next

Generate the client certificate for the agent. See [Regenerate an Agent Client Certificate](#).

Install the Agent on a Windows Platform Using the Windows Installer

You can install the End Point Operations Management agent on a Windows platform using a Windows installer.

You can perform a silent installation of the agent. See [Installing an End Point Operations Management Agent Silently on a Windows Machine](#).

Prerequisites

- Verify that you have sufficient privileges to deploy an End Point Operations Management agent. You must have vRealize Operations Manager user credentials that include a role that allows you to install End Point Operations Management agents. See [Roles and Privileges in vRealize Operations Manager](#).

- Verify that the installation directory for the End Point Operations Management agent does not contain a vRealize Hyperic agent installation.
- If you already have an End Point Operations Management agent installed on the machine, verify that it is not running.
- Verify that you do not have any End Point Operations Management or vRealize Hyperic agent installed on your environment before running the agent Windows installer.
- You must know the user name and password for the vRealize Operations Manager, the vRealize Operations Manager server address (FQDN), and the server certificate thumbprint value. You can see additional information about the certificate thumbprint in the procedure.

Procedure

- 1 Download the Windows installation EXE file that is appropriate for your Windows platform.

Operating System	RPM Bundle to Download
64bit Operating System	epops-agent-x86-64-win-version.exe
32bit Operating System	epops-agent-x86-win-version.exe

- 2 Double-click the file to open the installation wizard.
- 3 Complete the steps in the installation wizard.

Verify that the user and system locales are identical, and that the installation path contains only characters that are part of the system locale's code page. You can set user and system locales in the Regional Options or Regional Settings control panel.

Note the following information related to defining the server certificate thumbprint.

- The server certificate thumbprint is required to run a silent installation.
 - Either the SHA1 or SHA256 algorithm can be used for the thumbprint.
 - By default, the vRealize Operations Manager server generates a self-signed CA certificate that is used to sign the certificate of all the nodes in the cluster. In this case, the thumbprint must be the thumbprint of the CA certificate, to allow for the agent to communicate with all nodes.
 - As a vRealize Operations Manager administrator, you can import a custom certificate instead of using the default. In this instance, you must specify a thumbprint corresponding to that certificate as the value of this property.
 - To view the certificate thumbprint value, log into the vRealize Operations Manager Administration interface at <https://IP Address/admin> and click the **SSL Certificate** icon located on the right of the menu bar. Unless you replaced the original certificate with a custom certificate, the second thumbprint in the list is the correct one. If you did upload a custom certificate, the first thumbprint in the list is the correct one.
- 4 (Optional) Run `ep-agent.bat query` to verify if the agent is installed and running.

Results

The agent begins running on the Windows platform.

Caution The agent will run even if some of the parameters that you provided in the installation wizard are missing or invalid. Check the `wrapper.log` and `agent.log` files in the *product installation path*/log directory to verify that there are no installation errors.

Installing an End Point Operations Management Agent Silently on a Windows Machine

You can install an End Point Operations Management agent on a Windows machine using silent or very silent installation.

Silent and very silent installations are performed from a command line interface using a setup installer executable file.

Verify that you do not have any End Point Operations Management or vRealize Hyperic agent installed on your environment before running the agent Windows installer.

Use the following parameters to set up the installation process. For more information about these parameters, see [Specify the End Point Operations Management Agent Setup Properties](#).

Caution The parameters that you specify for the Windows installer are passed to the agent configuration without validation. If you provide an incorrect IP address or user credentials, the End Point Operations Management agent cannot start.

Table 1-2. Silent Command Line Installer Parameters

Parameter	Value	Mandatory /Optional	Comments
-serverAddress	FQDN/IP address	Mandatory	FQDN or IP address of the vRealize Operations Manager server.
-username	string	Mandatory	
-securePort	number	Optional	Default is 443
-password	string	Mandatory	
-serverCertificateThumbprint	string	Mandatory	The vRealize Operations Manager server certificate thumbprint. You must enclose the certificate thumbprint in opening and closing quotation marks, for example, -serverCertificateThumbprint "31:32:FA:1F:FD:78:1E:D8:9A:15:32:85:D7:FE:54:49:0A:1D:9F:6D" .

Parameters are available to define various other attributes for the installation process.

Table 1-3. Additional Silent Command Line Installer Parameters

Parameter	Default Value	Comments
/DIR	C:\ep-agent	Specifies the installation path. You cannot use spaces in the installation path, and you must connect the /DIR command and the installation path with an equal sign, for example, /DIR=C:\ep-agent.
/SILENT	none	Specifies that the installation is to be silent. In a silent installation, only the progress window appears.
/VERYSILENT	none	Specifies that the installation is to be very silent. In a very silent installation, the progress window does not appear, however installation error messages are displayed, as is the startup prompt if you did not disable it.

Install the Agent on an AIX Platform

You can install the End Point Operations Management agent on an AIX platform.

Prerequisites

- 1 Install IBM Java 7.
- 2 Add the latest JCE from the IBM JRE security directory: JAVA_INSTALLATION_DIR/jre/lib/security. For more information, see [Downloading and installing the unrestricted JCE policy files](#)

Procedure

- 1 When you configure the PATH variable, add /usr/java7_64/jre/bin:/usr/java7_64/bin or PATH=/usr/java7_64/jre/bin:/usr/java7_64/bin:\$PATH.
- 2 Configure HQ_JAVA_HOME=path_to_current_java_directory.
For more information on setting up and checking your AIX environment, see https://www.ibm.com/support/knowledgecenter/SSYKE2_7.0.0/com.ibm.java.aix.70.doc/diag/problem_determination/aix_setup.html.
- 3 Download the noJre version of the End Point Operations Management agent and install the agent on an AIX machine.
- 4 For agent installation information, see [Install the Agent on a Linux Platform from an Archive](#)

Install the Agent on a Solaris Platform

You can install the End Point Operations Management agent on a Solaris platform.

Prerequisites

- 1 Install Java 7 or above for Solaris from the Oracle site: https://java.com/en/download/help/solaris_install.xml

- 2 Add the latest JCE from <http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>

Procedure

- 1 When you configure the PATH variable, add /usr/java7_64/jre/bin:/usr/java7_64/bin or `PATH=/usr/java7_64/jre/bin:/usr/java7_64/bin:$PATH`.
- 2 Configure `HQ_JAVA_HOME=path_to_current_java_directory`.
- 3 Download and install the noJre version of the End Point Operations Management agent on a Solaris machine.
- 4 For agent installation information, see [Install the Agent on a Linux Platform from an Archive](#)

Java Prerequisites for the End Point Operations Management Agent

All End Point Operations Management agents require Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction policy files be included as part of the Java package.

Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction policy files are included in the JRE End Point Operations Management agent installation options.

You can install an End Point Operations Management agent package that does not contain JRE files, or choose to add JRE later.

If you select a non-JRE installation option, you must ensure that your Java package includes Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction policy files to enable registration of the End Point Operations Management agent. If you select a non-JRE option and your Java package does not include Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction policy files, you receive these error messages `Server might be down (or wrong IP/port were used)` and `Cannot support TLS_RSA_WITH_AES_256_CBC_SHA with currently installed providers`.

Configuring JRE Locations for End Point Operations Management Components

End Point Operations Management agents require a JRE. The platform-specific End Point Operations Management agent installers include a JRE. Platform-independent End Point Operations Management agent installers do not include a JRE.

If you select a non-JRE installation option, you must ensure that your Java package includes Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction policy files to enable registration of the End Point Operations Management agent. For more information, see [Java Prerequisites for the End Point Operations Management Agent](#).

Depending on your environment and the installation package that you use, you might need to define the location of the JRE for your agents. The following environments require JRE location configuration.

- Platform-specific agent installation on a machine that has its own JRE that you want to use.
- Platform-independent agent installation.

How the Agent Resolves its JRE

The agent resolves its JRE based on platform type.

UNIX-like Platforms

On UNIX-like platforms, the agent determines which JRE to use in the following order:

- 1 HQ_JAVA_HOME environment variable
- 2 Embedded JRE
- 3 JAVA_HOME environment variable

Linux Platforms

On Linux platforms, you use `export HQ_JAVA_HOME= path_to_current_java_directory` to define a system variable.

Windows Platforms

On Windows platforms, the agent resolves the JRE to use in the following order:

- 1 HQ_JAVA_HOME environment variable

The path defined in the variable must not contain spaces. Consider using a shortened version of the path, using the tild (~) character. For example, `c:\Program Files\Java\jre7` can become `c:\Progra~1\Java\jre7`. The number after the tild depends on the alphabetical order (where a = 1, b =2, and so on) of files whose name begins with `progra` in that directory.

- 2 Embedded JRE

You define a system variable from the **My Computer** menu. Select **Properties > Advanced > Environment Variables > System Variables > New**.

Because of a known issue with Windows, on Windows Server 2008 R2 and 2012 R2, Windows services might keep old values of system variables, even though they have been updated or removed. As a result, updates or removal of the HQ_JAVA_HOME system variable might not be propagated to the End Point Operations Management Agent service. In this event, the End Point Operations Management agent might use an obsolete value for HQ_JAVA_HOME, which causes it to use the wrong JRE version.

System Prerequisites for the End Point Operations Management Agent

If you do not define `localhost` as the loopback address, the End Point Operations Management agent does not register and the following error appears: `Connection failed. Server may be down (or wrong IP/port were used). Waiting for 10 seconds before retrying.`

As a workaround, complete the following steps:

Procedure

- 1 Open the hosts file `/etc/hosts` on Linux or `C:\Windows\System32\Drivers\etc\hosts` on Windows.
- 2 Modify the file to include a `localhost` mapping to the IPv4 `127.0.0.1` loopback address, using `127.0.0.1 localhost`.
- 3 Save the file.

Configure the End Point Operations Management Agent to vRealize Operations Manager Server Communication Properties

Before first agent startup, you can define the properties that enable the agent to communicate with the vRealize Operations Manager server, and other agent properties, in the `agent.properties` file of an agent. When you configure the agent in the properties file you can streamline the deployment for multiple agents.

If a properties file exists, back it up before you make configuration changes. If the agent does not have a properties file, create one.

An agent looks for its properties file in `AgentHome/conf`. This is the default location of `agent.properties`.

If the agent does not find the required properties for establishing communications with the vRealize Operations Manager server in either of these locations, it prompts for the property values at initial start up of the agent.

A number of steps are required to complete the configuration.

You can define some agent properties before or after the initial startup. You must always configure properties that control the following behaviors before initial startup.

- When the agent must use an SSL keystore that you manage, rather than a keystore that vRealize Operations Manager generates.
- When the agent must connect to the vRealize Operations Manager server through a proxy server.

Prerequisites

Verify that the vRealize Operations Manager server is running.

Procedure

- 1 [Activate End Point Operations Management Agent to vRealize Operations Manager Server Setup Properties](#)

In the `agent.properties` file, properties relating to communication between the End Point Operations Management agent and the vRealize Operations Manager server are inactive by default. You must activate them.

2 Specify the End Point Operations Management Agent Setup Properties

The `agent.properties` file contains properties that you can configure to manage communication.

3 Configure an End Point Operations Management Agent Keystore

The agent uses a self-signed certificate for internal communication, and a second certificate that is signed by the server during the agent registration process. By default, the certificates are stored in a keystore that is generated in the `data` folder. You can configure your own keystore for the agent to use.

4 Configure the End Point Operations Management Agent by Using the Configuration Dialog Box

The End Point Operations Management agent configuration dialog box appears in the shell when you start an agent that does not have configuration values that specify the location of the vRealize Operations Manager server. The dialog box prompts you to provide the address and port of the vRealize Operations Manager server, and other connection-related data.

5 Overriding Agent Configuration Properties

You can specify that vRealize Operations Manager override default agent properties when they differ from custom properties that you have defined.

6 End Point Operations Management Agent Properties

Multiple properties are supported in the `agent.properties` file for an End Point Operations Management agent. Not all supported properties are included by default in the `agent.properties` file.

What to do next

Start the End Point Operations Management agent.

Activate End Point Operations Management Agent to vRealize Operations Manager Server Setup Properties

In the `agent.properties` file, properties relating to communication between the End Point Operations Management agent and the vRealize Operations Manager server are inactive by default. You must activate them.

Procedure

- 1 In the `agent.properties` file, locate the following section.

```
## Use the following to automate agent setup
## using these properties.
##
## If any properties do not have values specified, the setup
## process prompts for their values.
##
## If the value to use during automatic setup is the default, use the string *default* as the
## value for the option.
```

- 2 Remove the hash tag at the beginning of each line to activate the properties.

```
#agent.setup.serverIP=localhost
#agent.setup.serverSSLPort=443
#agent.setup.serverLogin=username
#agent.setup.serverPword=password
```

The first time that you start the End Point Operations Management agent, if `agent.setup.serverPword` is inactive, and has a plain text value, the agent encrypts the value.

- 3 (Optional) Remove the hash tag at the beginning of the line `#agent.setup.serverCertificateThumbprint=` and provide a thumbprint value to activate pre-approval of the server certificate.

Specify the End Point Operations Management Agent Setup Properties

The `agent.properties` file contains properties that you can configure to manage communication.

Agent-server setup requires a minimum set of properties.

Procedure

- 1 Specify the location and credentials the agent must use to contact the vRealize Operations Manager server.

Property	Property Definition
agent.setup.serverIP	Specify the address or hostname of the vRealize Operations Manager server.
agent.setup.serverSSLPort	The default value is the standard SSL vRealize Operations Manager server listen port. If your server is configured for a different listen port, specify the port number.
agent.setup.serverLogin	Specify the user name for the agent to use when connecting to the vRealize Operations Manager server. If you change the value from the <code>username</code> default value, verify that the user account is correctly configured on the vRealize Operations Manager server.
agent.setup.serverPword	Specify the password for the agent to use, together with the vRealize Operations Manager user name, when connecting to the vRealize Operations Manager server. Verify that the password is the one configured in vRealize Operations Manager for the user account.

2 (Optional) Specify the vRealize Operations Manager server certificate thumbprint.

Property	Property Definition
agent.setup.serverCertificateThumbprint	<p>Provides details about the server certificate to trust.</p> <p>This parameter is required to run a silent installation.</p> <p>Either the SHA1 or SHA256 algorithm can be used for the thumbprint.</p> <p>By default, the vRealize Operations Manager server generates a self-signed CA certificate that is used to sign the certificate of all the nodes in the cluster. In this case, the thumbprint must be the thumbprint of the CA certificate, to allow for the agent to communicate with all nodes.</p> <p>As a vRealize Operations Manager administrator, you can import a custom certificate instead of using the default. In this instance, you must specify a thumbprint corresponding to that certificate as the value of this property.</p> <p>To view the certificate thumbprint value, log into the vRealize Operations Manager Administration interface at https://IP Address/admin and click the SSL Certificate icon located on the right of the menu bar. Unless you replaced the original certificate with a custom certificate, the second thumbprint in the list is the correct one. If you did upload a custom certificate, the first thumbprint in the list is the correct one.</p>

3 (Optional) Specify the location and file name of the platform token file.

This file is created by the agent during installation and contains the identity token for the platform object.

Property	Property Definition
Windows: agent.setup.tokenFileWindows	<p>Provides details about the location and name of the platform token file.</p> <p>The value cannot include backslash (\) or percentage(%) characters, or environment variables.</p>
Linux: agent.setup.tokenFileLinux	<p>Ensure that you use forward slashes (/) when specifying the Windows path.</p>

4 (Optional) Specify any other required properties by running the appropriate command.

Operating System	Command
Linux	<code>./bin/ep-agent.sh set-property <i>PropertyKey</i> <i>PropertyValue</i></code>
Windows	<code>./bin/ep-agent.bat set-property <i>PropertyKey</i> <i>PropertyValue</i></code>

The properties are encrypted in the `agent.properties` file.

Configure an End Point Operations Management Agent Keystore

The agent uses a self-signed certificate for internal communication, and a second certificate that is signed by the server during the agent registration process. By default, the certificates are stored in a keystore that is generated in the `data` folder. You can configure your own keystore for the agent to use.

Important To use your own keystore, you must perform this task before the first agent activation.

Procedure

- 1 In the `agent.properties` file, activate the `# agent.keystore.path=` and `# agent.keystore.password=` properties.

Define the full path to the keystore with `agent.keystore.path` and the keystore password with `agent.keystore.password`.

- 2 Add the `[agent.keystore.alias]` property to the properties file, and set it to the alias of the primary certificate or private key entry of the keystore primary certificate.

Configure the End Point Operations Management Agent by Using the Configuration Dialog Box

The End Point Operations Management agent configuration dialog box appears in the shell when you start an agent that does not have configuration values that specify the location of the vRealize Operations Manager server. The dialog box prompts you to provide the address and port of the vRealize Operations Manager server, and other connection-related data.

The agent configuration dialog box appears in these cases:

- The first time that you start an agent, if you did not supply one or more of the relevant properties in the `agent.properties` file.
- When you start an agent for which saved server connection data is corrupt or was removed.

You can also run the agent launcher to rerun the configuration dialog box.

Prerequisites

Verify that the server is running.

Procedure

- 1 Open a terminal window on the platform on which the agent is installed.
- 2 Navigate to the `AgentHome/bin` directory.
- 3 Run the agent launcher using the `start` or `setup` option.

Platform	Command
UNIX-like	<code>ep-agent.sh start</code>
Windows	<p>Install the Windows service for the agent, then run the it: <code>ep-agent.bat install ep-agent.bat start</code> command.</p> <p>When you configure an End Point Operations Management agent as a Windows service, make sure that the credentials that you specify are sufficient for the service to connect to the monitored technology. For example, if you have an End Point Operations Management agent that is running on Microsoft SQL Server, and only a specific user can log in to that server, the Windows service login must also be for that specific user.</p>

4 Respond to the prompts, noting the following as you move through the process.

Prompt	Description
Enter the server hostname or IP address	If the server is on the same machine as the agent, you can enter <code>localhost</code> . If a firewall is blocking traffic from the agent to the server, specify the address of the firewall.
Enter the server SSL port	Specify the SSL port on the vRealize Operations Manager server to which the agent must connect. The default port is 443.
The server has presented an untrusted certificate	If this warning appears, but your server is signed by a trusted certificate or you have updated the <code>thumbprint</code> property to contain the thumbprint, this agent might be subject to a man-in-the-middle attack. Review the displayed certificate thumbprint details carefully.
Enter your server username	Enter the name of a vRealize Operations Manager user with <code>agentManager</code> permissions.
Enter your server password	Enter the password for the specified vRealize Operations Manager. Do not store the password in the <code>agent.properties</code> file.

Results

The agent initiates a connection to the vRealize Operations Manager server and the server verifies that the agent is authenticated to communicate with it.

The server generates a client certificate that includes the agent token. The message `The agent has been successfully registered` appears. The agent starts discovering the platform and supported products running on it.

Overriding Agent Configuration Properties

You can specify that vRealize Operations Manager override default agent properties when they differ from custom properties that you have defined.

In the Advanced section of the Edit Object dialog box, if you set the **Override agent configuration data** to `false`, default agent configuration data is applied. If you set **Override agent configuration data** to `true`, the default agent parameter values are ignored if you have set alternative values, and the values that you set are applied.

If you set the value of **Override agent configuration data** to `true` when editing an MSSQL object (MSSQL, MSSQL Database, MSSQL Reporting Services, MSSQL Analysis Service, or MSSQL Agent) that runs in a cluster, it might result in inconsistent behavior.

End Point Operations Management Agent Properties

Multiple properties are supported in the `agent.properties` file for an End Point Operations Management agent. Not all supported properties are included by default in the `agent.properties` file.

You must add any properties that you want to use that are not included in the default `agent.properties` file.

You can encrypt properties in the `agent.properties` file to enable silent installation.

Encrypt End Point Operations Management Agent Property Values

After you have installed an End Point Operations Management agent, you can use it to add encrypted values to the `agent.properties` file to enable silent installation.

For example, to specify the user password, you can run `./bin/ep-agent.sh set-property agent.setup.serverPword serverPasswordValue` to add the following line to the `agent.properties` file.

```
agent.setup.serverPword = ENC(4FyUf6m/c5i+RriaNpSEQ1WKGb4y
+Dhp7213XQiyvtwI4tM1bGJfZMBPG23KnsUWu30KrW35gB+Ms20snM4TDg==)
```

The key that was used to encrypt the value is saved in `AgentHome/conf/agent.scu`. If you encrypt other values, the key that was used to encrypt the first value is used.

Prerequisites

Verify that the End Point Operations Management agent can access `AgentHome/conf/agent.scu`. Following the encryption of any agent-to-server connection properties, the agent must be able to access this file to start.

Procedure

- ◆ Open a command prompt and run `./bin/ep-agent.sh set-property agent.setup.propertyName propertyValue`.

Results

The key that was used to encrypt the value is saved in `AgentHome/conf/agent.scu`.

What to do next

If your agent deployment strategy involves distributing a standard `agent.properties` file to all agents, you must also distribute `agent.scu`. See [Install Multiple End Point Operations Management Agents Simultaneously](#).

Adding Properties to the agent.properties File

You must add any properties that you want to use that are not included in the default `agent.properties` file.

Following is a list of the available properties.

- [agent.keystore.alias Property](#)
This property configures the name of the user-managed keystore for the agent for agents configured for unidirectional communication with the vRealize Operations Manager server.
- [agent.keystore.password Property](#)
This property configures the password for an End Point Operations Management agent's SSL keystore.
- [agent.keystore.path Property](#)
This property configures the location of a End Point Operations Management agent's SSL keystore.

- [agent.listenPort Property](#)

This property specifies the port where the End Point Operations Management agent listens to receive communication from the vRealize Operations Manager server.

- [agent.logDir Property](#)

You can add this property to the `agent.properties` file to specify the directory where the End Point Operations Management agent writes its log file. If you do not specify a fully qualified path, `agent.logDir` is evaluated relative to the agent installation directory.

- [agent.logFile Property](#)

The path and name of the agent log file.

- [agent.logLevel Property](#)

The level of detail of the messages the agent writes to the log file.

- [agent.logLevel.SystemErr Property](#)

Redirects `System.err` to the `agent.log` file.

- [agent.logLevel.SystemOut Property](#)

Redirects `System.out` to the `agent.log` file.

- [agent.proxyHost Property](#)

The host name or IP address of the proxy server that the End Point Operations Management agent must connect to first when establishing a connection to the vRealize Operations Manager server.

- [agent.proxyPort Property](#)

The port number of the proxy server that the End Point Operations Management agent must connect to first when establishing a connection to the vRealize Operations Manager server.

- [agent.setup.acceptUnverifiedCertificate Property](#)

This property controls whether an End Point Operations Management agent issues a warning when the vRealize Operations Manager server presents an SSL certificate that is not in the agent's keystore, and is either self-signed or signed by a different certificate authority than the one that signed the agent's SSL certificate.

- [agent.setup.camIP Property](#)

Use this property to define the IP address of the vRealize Operations Manager server for the agent. The End Point Operations Management agent reads this value only in the event that it cannot find connection configuration in its data directory.

- [agent.setup.camLogin Property](#)

At first startup after installation, use this property to define the End Point Operations Management agent user name to use when the agent is registering itself with the server.

- [agent.setup.camPort Property](#)

At first startup after installation, use this property to define the End Point Operations Management agent server port to use for non-secure communications with the server.

- [agent.setup.camPword Property](#)

Use this property to define the password that the End Point Operations Management agent uses when connecting to the vRealize Operations Manager server, so that the agent does not prompt a user to supply the password interactively at first startup.

- [agent.setup.camSecure](#)

This property is used when you are registering the End Point Operations Management with the vRealize Operations Manager server to communicate using encryption.

- [agent.setup.camSSLPort Property](#)

At first startup after installation, use this property to define the End Point Operations Management agent server port to use for SSL communications with the server.

- [agent.setup.resetupToken Property](#)

Use this property to configure an End Point Operations Management agent to create a new token to use for authentication with the server at startup. Regenerating a token is useful if the agent cannot connect to the server because the token has been deleted or corrupted.

- [agent.setup.unidirectional Property](#)

Enables unidirectional communications between the End Point Operations Management agent and vRealize Operations Manager server.

- [agent.startupTimeOut Property](#)

The number of seconds that the End Point Operations Management agent startup script waits before determining that the agent has not started up successfully. If the agent is found to not be listening for requests within this period, an error is logged, and the startup script times out.

- [autoinventory.defaultScan.interval.millis Property](#)

Specifies how frequently the End Point Operations Management agent performs a default autoinventory scan.

- [autoinventory.runtimeScan.interval.millis Property](#)

Specifies how frequently an End Point Operations Management agent performs a runtime scan.

- [http.useragent Property](#)

Defines the value for the user-agent request header in HTTP requests issued by the End Point Operations Management agent.

- [log4j Properties](#)

The log4j properties for the End Point Operations Management agent are described here.

- [platform.log_track.eventfmt Property](#)

Specifies the content and format of the Windows event attributes that an End Point Operations Management agent includes when logging a Windows event as an event in vRealize Operations Manager.

- [plugins.exclude Property](#)

Specifies plug-ins that the End Point Operations Management agent does not load at startup. This is useful for reducing an agent's memory footprint.

- [plugins.include Property](#)

Specifies plug-ins that the End Point Operations Management agent loads at startup. This is useful for reducing the agent's memory footprint.

- [postgresql.database.name.format Property](#)

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Database and vPostgreSQL Database database types.

- [postgresql.index.name.format Property](#)

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Index and vPostgreSQL Index index types.

- [postgresql.server.name.format Property](#)

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL and vPostgreSQL server types.

- [postgresql.table.name.format Property](#)

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Table and vPostgreSQL Table table types.

- [scheduleThread.cancelTimeout Property](#)

This property specifies the maximum time, in milliseconds, that the ScheduleThread allows a metric collection process to run before attempting to interrupt it.

- [scheduleThread.fetchLogTimeout Property](#)

This property controls when a warning message is issued for a long-running metric collection process.

- [scheduleThread.poolsize Property](#)

This property enables a plug-in to use multiple threads for metric collection. The property can increase metric throughput for plug-ins known to be thread-safe.

- [scheduleThread.queueSize Property](#)

Use this property to limit the metric collection queue size (the number of metrics) for a plug-in.

- [sigar.mirror.procnet Property](#)

mirror /proc/net/tcp on Linux.

- [sigar.pdh.enableTranslation Property](#)

Use this property to enable translation based on the detected locale of the operating system.

■ [snmpTrapReceiver.listenAddress Property](#)

Specifies the port on which the End Point Operations Management agent listens for SNMP traps

agent.keystore.alias Property

This property configures the name of the user-managed keystore for the agent for agents configured for unidirectional communication with the vRealize Operations Manager server.

Example: Defining the Name of a Keystore

Given this user-managed keystore for a unidirectional agent

```
hq self-signed cert), Jul 27, 2011, trustedCertEntry,
Certificate fingerprint (MD5): 98:FF:B8:3D:25:74:23:68:6A:CB:0B:9C:20:88:74:CE
hq-agent, Jul 27, 2011, PrivateKeyEntry,
Certificate fingerprint (MD5): 03:09:C4:BC:20:9E:9A:32:DC:B2:E8:29:C0:3C:FE:38
```

you define the name of the keystore like this

```
agent.keystore.alias=hq-agent
```

If the value of this property does not match the keystore name, agent-server communication fails.

Default

The default behavior of the agent is to look for the hq keystore.

For unidirectional agents with user-managed keystores, you must define the keystore name using this property.

agent.keystore.password Property

This property configures the password for an End Point Operations Management agent's SSL keystore.

Define the location of the keystore using the [agent.keystore.path Property](#) property.

By default, the first time you start the End Point Operations Management agent following installation, if `agent.keystore.password` is uncommented and has a plain text value, the agent automatically encrypts the property value. You can encrypt this property value yourself, prior to starting the agent.

It is good practice to specify the same password for the agent keystore as for the agent private key.

Default

By default, the `agent.properties` file does not include this property.

agent.keystore.path Property

This property configures the location of a End Point Operations Management agent's SSL keystore.

Specify the full path to the keystore. Define the password for the keystore using the `agent.keystore.password` property. See [agent.keystore.password Property](#).

Specifying the Keystore Path on Windows

On Windows platforms, specify the path to the keystore in this format.

```
C:/Documents and Settings/Desktop/keystore
```

Default

`AgentHome/data/keystore.`

`agent.listenPort` Property

This property specifies the port where the End Point Operations Management agent listens to receive communication from the vRealize Operations Manager server.

The property is not required for unidirectional communication.

`agent.logDir` Property

You can add this property to the `agent.properties` file to specify the directory where the End Point Operations Management agent writes its log file. If you do not specify a fully qualified path, `agent.logDir` is evaluated relative to the agent installation directory.

To change the location for the agent log file, enter a path relative to the agent installation directory, or a fully qualified path.

Note that the name of the agent log file is configured with the `agent.logFile` property.

Default

By default, the `agent.properties` file does not include this property.

The default behavior is `agent.logDir=log`, resulting in the agent log file being written to the `AgentHome/log` directory.

`agent.logFile` Property

The path and name of the agent log file.

Default

In the `agent.properties` file, the default setting for the `agent.LogFile` property is made up of a variable and a string

```
agent.logFile=${agent.logDir}\agent.log
```

where

- `agent.logDir` is a variable that supplies the value of an identically named agent property. By default, the value of `agent.logDir` is `log`, interpreted relative to the agent installation directory.
- `agent.log` is the name for the agent log file.

By default, the agent log file is named `agent.log`, and is written to the `AgentHome/log` directory.

`agent.logLevel` Property

The level of detail of the messages the agent writes to the log file.

Permitted values are INFO and DEBUG.

Default

INFO

agent.logLevel.SystemErr Property

Redirects System.err to the agent.log file.

Commenting out this setting causes System.err to be directed to agent.log.startup.

Default

ERROR

agent.logLevel.SystemOut Property

Redirects System.out to the agent.log file.

Commenting out this setting causes System.out to be directed to agent.log.startup.

Default

INFO

agent.proxyHost Property

The host name or IP address of the proxy server that the End Point Operations Management agent must connect to first when establishing a connection to the vRealize Operations Manager server.

This property is supported for agents configured for unidirectional communication.

Use this property in conjunction with agent.proxyPort and agent.setup.unidirectional.

Default

None

agent.proxyPort Property

The port number of the proxy server that the End Point Operations Management agent must connect to first when establishing a connection to the vRealize Operations Manager server.

This property is supported for agents configured for unidirectional communication.

Use this property in conjunction with agent.proxyPort and agent.setup.unidirectional.

Default

None

agent.setup.acceptUnverifiedCertificate Property

This property controls whether an End Point Operations Management agent issues a warning when the vRealize Operations Manager server presents an SSL certificate that is not in the agent's keystore, and is either self-signed or signed by a different certificate authority than the one that signed the agent's SSL certificate.

When the default is used, the agent issues the warning

```
The authenticity of host 'localhost' can't be established.
Are you sure you want to continue connecting? [default=no]:
```

If you respond **yes**, the agent imports the server's certificate and will continue to trust the certificate from this point on.

Default

`agent.setup.acceptUnverifiedCertificate=no`

`agent.setup.camIP` Property

Use this property to define the IP address of the vRealize Operations Manager server for the agent. The End Point Operations Management agent reads this value only in the event that it cannot find connection configuration in its data directory.

You can specify this and other `agent.setup.*` properties to reduce the user interaction required to configure an agent to communicate with the server.

The value can be provided as an IP address or a fully qualified domain name. To identify an server on the same host as the server, set the value to 127.0.0.1.

If there is a firewall between the agent and server, specify the address of the firewall, and configure the firewall to forward traffic on port 7080, or 7443 if you use the SSL port, to the vRealize Operations Manager server.

Default

Commented out, localhost.

`agent.setup.camLogin` Property

At first startup after installation, use this property to define the End Point Operations Management agent user name to use when the agent is registering itself with the server.

The permission required on the server for this initialization is Create, for platforms.

Log in from the agent to the server is only required during the initial configuration of the agent.

The agent reads this value only in the event that it cannot find connection configuration in its data directory.

You can specify this and other `agent.setup.*` properties to reduce the user interaction required to configure an agent to communicate with the server.

Default

Commented out, hqadmin.

`agent.setup.camPort` Property

At first startup after installation, use this property to define the End Point Operations Management agent server port to use for non-secure communications with the server.

The agent reads this value only in the event that it cannot find connection configuration in its data directory.

You can specify this and other `agent.setup.*` properties to reduce the user interaction required to configure an agent to communicate with the server.

Default

Commented out 7080.

agent.setup.camPword Property

Use this property to define the password that the End Point Operations Management agent uses when connecting to the vRealize Operations Manager server, so that the agent does not prompt a user to supply the password interactively at first startup.

The password for the user is that specified by `agent.setup.camLogin`.

The agent reads this value only in the event that it cannot find connection configuration in its data directory.

You can specify this and other `agent.setup.*` properties to reduce the user interaction required to configure an agent to communicate with the server.

The first time you start the End Point Operations Management agent after installation, if `agent.keystore.password` is uncommented and has a plain text value, the agent automatically encrypts the property value. You can encrypt these property values prior to starting the agent.

Default

Commented out `hqadmin`.

agent.setup.camSecure

This property is used when you are registering the End Point Operations Management with the vRealize Operations Manager server to communicate using encryption.

Use `yes=secure`, `encrypted`, or `SSL`, as appropriate, to encrypt communication.

Use `no=unencrypted` for unencrypted communication.

agent.setup.camSSLPort Property

At first startup after installation, use this property to define the End Point Operations Management agent server port to use for SSL communications with the server.

The agent reads this value only in the event that it cannot find connection configuration in its data directory.

You can specify this and other `agent.setup.*` properties to reduce the user interaction required to configure an agent to communicate with the server.

Default

Commented out `7443`.

agent.setup.resetupToken Property

Use this property to configure an End Point Operations Management agent to create a new token to use for authentication with the server at startup. Regenerating a token is useful if the agent cannot connect to the server because the token has been deleted or corrupted.

The agent reads this value only in the event that it cannot find connection configuration in its data directory.

Regardless of the value of this property, an agent generates a token the first time it is started after installation.

Default

Commented out no.

agent.setup.unidirectional Property

Enables unidirectional communications between the End Point Operations Management agent and vRealize Operations Manager server.

If you configure an agent for unidirectional communication, all communication with the server is initiated by the agent.

For a unidirectional agent with a user-managed keystore, you must configure the keystore name in the `agent.properties` file.

Default

Commented out no.

agent.startupTimeout Property

The number of seconds that the End Point Operations Management agent startup script waits before determining that the agent has not started up successfully. If the agent is found to not be listening for requests within this period, an error is logged, and the startup script times out.

Default

By default, the `agent.properties` file does not include this property.

The default behavior of the agent is to timeout after 300 seconds.

autoinventory.defaultScan.interval.millis Property

Specifies how frequently the End Point Operations Management agent performs a default autoinventory scan.

The default scan detects server and platform services objects, typically using the process table or the Windows registry. Default scans are less resource-intensive than runtime scans.

Default

The agent performs the default scan at startup and every 15 minutes thereafter.

Commented out 86,400,000 milliseconds, or one day.

autoinventory.runtimeScan.interval.millis Property

Specifies how frequently an End Point Operations Management agent performs a runtime scan.

A runtime scan may use more resource-intensive methods to detect services than a default scan. For example, a runtime scan might involve issuing an SQL query or looking up an MBean.

Default

86,400,000 milliseconds, or one day.

http.useragent Property

Defines the value for the user-agent request header in HTTP requests issued by the End Point Operations Management agent.

You can use `http.useragent` to define a user-agent value that is consistent across upgrades.

By default, the `agent.properties` file does not include this property.

Default

By default, the user-agent in agent requests includes the End Point Operations Management agent version, so changes when the agent is upgraded. If a target HTTP server is configured to block requests with an unknown user-agent, agent requests fail after an agent upgrade.

Hyperic-HQ-Agent/Version, for example, Hyperic-HQ-Agent/4.1.2-EE.

log4j Properties

The log4j properties for the End Point Operations Management agent are described here.

```
log4j.rootLogger=${agent.logLevel}, R

log4j.appender.R.File=${agent.logFile}
log4j.appender.R.MaxBackupIndex=1
log4j.appender.R.MaxFileSize=5000KB
log4j.appender.R.layout.ConversionPattern=%d{dd-MM-yyyy HH:mm:ss,SSS z} %-5p [%t] [%c{1}@%L] %m%n
log4j.appender.R.layout=org.apache.log4j.PatternLayout
log4j.appender.R=org.apache.log4j.RollingFileAppender

##
## Disable overly verbose logging
##
log4j.logger.org.apache.http=ERROR
log4j.logger.org.springframework.web.client.RestTemplate=ERROR
log4j.logger.org.hyperic.hq.measurement.agent.server.SenderThread=INFO
log4j.logger.org.hyperic.hq.agent.server.AgentDListProvider=INFO
log4j.logger.org.hyperic.hq.agent.server.MeasurementSchedule=INFO
log4j.logger.org.hyperic.util.units=INFO
log4j.logger.org.hyperic.hq.product.pluginxml=INFO

# Only log errors from naming context
log4j.category.org.jnp.interfaces.NamingContext=ERROR
log4j.category.org.apache.axis=ERROR

#Agent Subsystems: Uncomment individual subsystems to see debug messages.
#-----
#log4j.logger.org.hyperic.hq.autoinventory=DEBUG
#log4j.logger.org.hyperic.hq.livedata=DEBUG
#log4j.logger.org.hyperic.hq.measurement=DEBUG
#log4j.logger.org.hyperic.hq.control=DEBUG

#Agent Plugin Implementations
#log4j.logger.org.hyperic.hq.product=DEBUG

#Server Communication
#log4j.logger.org.hyperic.hq.bizapp.client.AgentCallbackClient=DEBUG

#Server Realtime commands dispatcher
#log4j.logger.org.hyperic.hq.agent.server.CommandDispatcher=DEBUG

#Agent Configuration parser
#log4j.logger.org.hyperic.hq.agent.AgentConfig=DEBUG

#Agent plugins loader
#log4j.logger.org.hyperic.util.PluginLoader=DEBUG
```

```
#Agent Metrics Scheduler (Scheduling tasks definitions & executions)
#log4j.logger.org.hyperic.hq.agent.server.session.AgentSynchronizer.SchedulerThread=DEBUG

#Agent Plugin Managers
#log4j.logger.org.hyperic.hq.product.MeasurementPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.AutoinventoryPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ConfigTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LogTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LiveDataPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ControlPluginManager=DEBUG
```

platform.log_track.eventfmt Property

Specifies the content and format of the Windows event attributes that an End Point Operations Management agent includes when logging a Windows event as an event in vRealize Operations Manager.

By default, the agent.properties file does not include this property.

Default

When Windows log tracking is enabled, an entry in the form [Timestamp] Log Message (EventLogName):EventLogName:EventAttributes is logged for events that match the criteria you specified on the resource's Configuration Properties page.

Attribute	Description
Timestamp	When the event occurred
Log Message	A text string
EventLogName	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, the log entry: 04/19/2010 06:06 AM Log Message (SYSTEM): SYSTEM: Print: Printer HP LaserJet 6P was paused. is for a 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.

Configuration

Use the following parameters 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.

Parameter	Description
%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.

For example, with the property setting `platform.log_track.eventfmt=%user%%computer% %source %:%event%:%message%`, the End Point Operations Management agent writes the following data when logging the Windows event 04/19/2010 06:06 AM Log Message (SYSTEM): SYSTEM:

HP_Administrator@Office Print:7:Printer HP LaserJet 6P was paused.. This entry is for a Windows event written to the Windows system event log at 6:06 AM on 04/19/2010. The software associated with the event was running as "HP_Administrator" on the host "Office". The Windows event's Source, Event, and Message attributes, are "Print", "7", and "Printer HP LaserJet 6P was paused.", respectively.

`plugins.exclude` Property

Specifies plug-ins that the End Point Operations Management agent does not load at startup. This is useful for reducing an agent's memory footprint.

Usage

Supply a comma-separated list of plug-ins to exclude. For example,

```
plugins.exclude=jboss,apache,mysql
```

`plugins.include` Property

Specifies plug-ins that the End Point Operations Management agent loads at startup. This is useful for reducing the agent's memory footprint.

Usage

Supply a comma-separated list of plug-ins to include. For example,

```
plugins.include=weblogic,apache
```

`postgresql.database.name.format` Property

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Database and vPostgreSQL Database database types.

By default, the name of a PostgreSQL or vPostgreSQL database is Database *DatabaseName*, where *DatabaseName* is the auto-discovered name of the database.

To use a different naming convention, define `postgresql.database.name.format`. The variable data you use must be available from the PostgreSQL plug-in.

Use the following syntax to specify the default table name assigned by the plug-in,

```
Database ${db}
```

where

`postgresql.db` is the auto-discovered name of the PostgreSQL or vPostgreSQL database.

Default

By default, the `agent.properties` file does not include this property.

`postgresql.index.name.format` Property

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Index and vPostgreSQL Index index types.

By default, the name of a PostgreSQL or vPostgreSQL index is `Index DatabaseName.Schema.Index`, comprising the following variables

Variable	Description
DatabaseName	The auto-discovered name of the database.
Schema	The auto-discovered schema for the database.
Index	The auto-discovered name of the index.

To use a different naming convention, define `postgresql.index.name.format`. The variable data you use must be available from the PostgreSQL plug-in.

Use the following syntax to specify the default index name assigned by the plug-in,

```
Index ${db}.${schema}.${index}
```

where

Attribute	Description
db	Identifies the platform that hosts the PostgreSQL or vPostgreSQL server.
schema	Identifies the schema associated with the table.
index	The index name in PostgreSQL.

Default

By default, the `agent.properties` file does not include this property.

postgresql.server.name.format Property

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL and vPostgreSQL server types.

By default, the name of a PostgreSQL or vPostgreSQL server is `Host:Port`, comprising the following variables

Variable	Description
Host	The FQDN of the platform that hosts the server.
Port	The PostgreSQL listen port.

To use a different naming convention, define `postgresql.server.name.format`. The variable data you use must be available from the PostgreSQL plug-in.

Use the following syntax to specify the default server name assigned by the plug-in,

```
${postgresql.host}:${postgresql.port}
```

where

Attribute	Description
postgresql.host	Identifies the FQDN of the hosting platform.
postgresql.port	Identifies the database listen port.

Default

By default, the `agent.properties` file does not include this property.

postgresql.table.name.format Property

This property specifies the format of the name that the PostgreSQL plug-in assigns to auto-discovered PostgreSQL Table and vPostgreSQL Table table types.

By default, the name of a PostgreSQL or vPostgreSQL table is Table *DatabaseName.Schema.Table*, comprising the following variables

Variable	Description
DatabaseName	The auto-discovered name of the database.
Schema	The auto-discovered schema for the database.
Table	The auto-discovered name of the table.

To use a different naming convention, define `postgresql.table.name.format`. The variable data you use must be available from the PostgreSQL plug-in.

Use the following syntax to specify the default table name assigned by the plug-in,

```
Table ${db}.${schema}.${table}
```

where

Attribute	Description
db	Identifies the platform that hosts the PostgreSQL or vPostgreSQL server.
schema	Identifies the schema associated with the table.
table	The table name in PostgreSQL.

Default

By default, the `agent.properties` file does not include this property.

scheduleThread.cancelTimeout Property

This property specifies the maximum time, in milliseconds, that the `ScheduleThread` allows a metric collection process to run before attempting to interrupt it.

When the timeout is exceeded, collection of the metric is interrupted, if it is in a `wait()`, `sleep()` or non-blocking `read()` state.

Usage

```
scheduleThread.cancelTimeout=5000
```

Default

5000 milliseconds.

scheduleThread.fetchLogTimeout Property

This property controls when a warning message is issued for a long-running metric collection process.

If a metric collection process exceeds the value of this property, which is measured in milliseconds, the agent writes a warning message to the `agent.log` file.

Usage

```
scheduleThread.fetchLogTimeout=2000
```

Default

2000 milliseconds.

scheduleThread.poolsize Property

This property enables a plug-in to use multiple threads for metric collection. The property can increase metric throughput for plug-ins known to be thread-safe.

Usage

Specify the plug-in by name and the number of threads to allocate for metric collection

```
scheduleThread.poolsize.PluginName=2
```

where *PluginName* is the name of the plug-in to which you are allocating threads. For example,

```
scheduleThread.poolsize.vsphere=2
```

Default

1

scheduleThread.queueSize Property

Use this property to limit the metric collection queue size (the number of metrics) for a plug-in.

Usage

Specify the plug-in by name and the maximum metric queue length number:

```
scheduleThread.queueSize.PluginName=15000
```

where *PluginName* is the name of the plug-in on which you are imposing a metric limit.

For example,

```
scheduleThread.queueSize.vsphere=15000
```

Default

1000

sigar.mirror.procnets Property

mirror /proc/net/tcp on Linux.

Default

true

sigar.pdh.enableTranslation Property

Use this property to enable translation based on the detected locale of the operating system.

snmpTrapReceiver.listenAddress Property

Specifies the port on which the End Point Operations Management agent listens for SNMP traps

By default, the `agent.properties` file does not include this property.

Typically SNMP uses the UDP port 162 for trap messages. This port is in the privileged range, so an agent listening for trap messages on it must run as root, or as an administrative user on Windows.

You can run the agent in the context of a non-administrative user, by configuring the agent to listen for trap messages on an unprivileged port.

Usage

Specify an IP address (or `0.0.0.0` to specify all interfaces on the platform) and the port for UDP communications in the format

```
snmpTrapReceiver.listenAddress=udp:IP_address/port
```

To enable the End Point Operations Management agent to receive SNMP traps on an unprivileged port, specify port 1024 or higher. The following setting allows the agent to receive traps on any interface on the platform, on UDP port 1620.

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

Managing Agent Registration on vRealize Operations Manager Servers

The End Point Operations Management agents identify themselves to the server using client certificates. The agent registration process generates the client certificate.

The client certificate includes a token that is used as the unique identifier. If you suspect that a client certificate was stolen or compromised, you must replace the certificate.

You must have AgentManager credentials to perform the agent registration process.

If you remove and reinstall an agent by removing the data directory, the agent token is retained to enable data continuity. See [Understanding Agent Uninstallation and Reinstallation Implications](#).

Regenerate an Agent Client Certificate

An End Point Operations Management agent client certificate might expire and need to be replaced. For example, you might replace a certificate that you suspected was corrupt or compromised.

Prerequisites

Verify that you have sufficient privileges to deploy an End Point Operations Management agent. You must have vRealize Operations Manager user credentials that include a role that allows you to install End Point Operations Management agents. See [Roles and Privileges in vRealize Operations Manager](#).

Procedure

- ◆ Start the registration process by running the `setup` command that is appropriate for the operating system on which the agent is running.

Operating System	Run Command
Linux	<code>ep-agent.sh setup</code>
Windows	<code>ep-agent.bat setup</code>

Results

The agent installer runs the `setup`, requests a new certificate from the server, and imports the new certificate to the keystore.

Securing Communications with the Server

Communication from an End Point Operations Management agent to the vRealize Operations Manager server is unidirectional, however both parties must be authenticated. Communication is always secured using transport layer security (TLS).

The first time an agent initiates a connection to the vRealize Operations Manager server following installation, the server presents its SSL certificate to the agent.

If the agent trusts the certificate that the server presented, the agent imports the server's certificate to its own keystore.

The agent trusts a server certificate if that certificate, or one of its issuers (CA) already exists in the agent's keystore.

By default, if the agent does not trust the certificate that the server presents, the agent issues a warning. You can choose to trust the certificate, or to terminate the configuration process. The vRealize Operations Manager server and the agent do not import untrusted certificates unless you respond yes to the warning prompt.

You can configure the agent to accept a specific thumb print without warning by specifying the thumb print of the certificate for the vRealize Operations Manager server.

By default, the vRealize Operations Manager server generates a self-signed CA certificate that is used to sign the certificate of all the nodes in the cluster. In this case, the thumbprint must be the thumbprint of the issuer, to allow for the agent to communicate with all nodes.

As a vRealize Operations Manager administrator, you can import a custom certificate instead of using the default. In this instance, you must specify a thumbprint corresponding to that certificate as the value of this property.

Either the SHA1 or SHA256 algorithm can be used for the thumbprint.

Launching Agents from a Command Line

You can launch agents from a command line on both Linux and Windows operating systems.

Use the appropriate process for your operating system.

If you are deleting the data directory, do not use Windows Services to stop and start an End Point Operations Management agent. Stop the agent using `epops-agent.bat stop`. Delete the data directory, then start the agent using `epops-agent.bat start`.

Run the Agent Launcher from a Linux Command Line

You can initiate the agent launcher and agent lifecycle commands with the `epops-agent.sh` script in the `AgentHome/bin` directory.

Procedure

- 1 Open a command shell or terminal window.
- 2 Enter the required command, using the format `sh epops-agent.sh command`, where *command* is one of the following.

Option	Description
start	Starts the agent as a daemon process.
stop	Stops the agent's JVM process.
restart	Stops and then starts the agent's JVM process.
status	Queries the status of the agent's JVM process.
dump	Runs a thread dump for the agent process, and writes the result to the <code>agent.log</code> file in <code>AgentHome/Log</code> .
ping	Pings the agent process.
setup	Re-registers the certificate using the existing token.

Run the Agent Launcher from a Windows Command Line

You can initiate the agent launcher and agent lifecycle commands with the `epops-agent.bat` script in the `AgentHome/bin` directory.

Procedure

- 1 Open a terminal window.
- 2 Enter the required command, using the format `epops-agent.bat command`, where *command* is one of the following.

Option	Description
install	Installs the agent NT service. You must run <code>start</code> after running <code>install</code> .
start	Starts the agent as an NT service.
stop	Stops the agent as an NT service.
remove	Removes the agent's service from the NT service table.
query	Queries the current status of the agent NT service (status).
dump	Runs a thread dump for the agent process, and writes the result to the <code>agent.log</code> file in <code>AgentHome/Log</code> .

Option	Description
ping	Pings the agent process.
setup	Re-registers the certificate using the existing token.

Managing an End Point Operations Management Agent on a Cloned Virtual Machine

When you clone a virtual machine that is running an End Point Operations Management agent that is collecting data, there are processes that you must complete related to data continuity to ensure data continuity.

Cloning a Virtual Machine to Delete the Original Virtual Machine

If you are cloning the virtual machine so that you can delete the original virtual machine, you need to verify that the original machine is deleted from the vCenter Server and from vRealize Operations Manager so that the new operating system to virtual machine relationship can be created.

Cloning a Virtual Machine to Run Independently of the Original Machine

If you are cloning the virtual machine so that you can run the two machines independently of the other, the cloned machine requires a new agent because an agent can only monitor a single machine.

Procedure

- ◆ On the cloned machine, delete the End Point Operations Management token and the data folder, according to the operating system of the machine.

Operating System	Process
Linux	Stop the End Point Operations Management services and delete the End Point Operations Management token and the data folder.
Windows	<ol style="list-style-type: none"> 1 Run <code>epops-agent remove</code>. 2 Remove the agent token and the data folder. 3 Run <code>epops-agent install</code>. 4 Run <code>epops-agent start</code>.

Moving Virtual Machines between vCenter Server Instances

When you move a virtual machine from one vCenter Server to another, vRealize Operations Manager preserves the unique object ID, identifiers, and historical data without creating any duplicate resources. This enables the new operating system to create a relationship with the migrated virtual machine.

Understanding Agent Uninstallation and Reinstallation Implications

When you uninstall or reinstall an End Point Operations Management agent, various elements are affected, including existing metrics that the agent has collected, and the identification token that enables a reinstalled agent to report on the previously discovered objects on the server. To ensure that you maintain data continuity, it is important that you are aware of the implications of uninstalling and reinstalling an agent.

There are two key locations related to the agent that are preserved when you uninstall an agent. Before reinstalling the agent, you must decide whether to retain or delete the files.

- The `/data` folder is created during agent installation. It contains the keystore, unless you chose a different location for it, and other data related to the currently installed agent.
- The `epops-token` platform token file is created before agent registration and is stored as follows:
 - Linux: `/etc/vmware/epops-token`
 - Windows: `%PROGRAMDATA%/VMware/EP Ops Agent/epops-token`

When you uninstall an agent, you must delete the `/data` folder. This does not affect data continuity.

However, to enable data continuity it is important that you do not delete the `epops-token` file. This file contains the identity token for the platform object. Following agent reinstallation, the token enables the agent to be synchronized with the previously discovered objects on the server.

When you reinstall the agent, the system notifies you whether it found an existing token, and provides its identifier. If a token is found, the system uses that token. If a token is not found, the system creates a new one. In the case of an error, the system prompts you to provide either a location and file name for the existing token file, or a location and file name for a new one.

The method that you use to uninstall an agent depends on how it was installed.

- [Uninstall an Agent that was Installed from an Archive](#)
You can use this procedure to uninstall agents that you installed on virtual machines in your environment from an archive.
- [Uninstall an Agent that was Installed Using an RPM Package](#)
You can use this procedure to uninstall agents that you installed on virtual machines in your environment using an RPM package.
- [Uninstall an Agent that was Installed Using a Windows Executable](#)
You can use this procedure to uninstall agents that you installed on virtual machines in your environment from a Windows EXE file.
- [Reinstall an Agent](#)
If you change the IP address, hostname or port number of the vRealize Operations Manager server, you need to uninstall and reinstall your agents.

Uninstall an Agent that was Installed from an Archive

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from an archive.

Prerequisites

Verify that the agent is stopped.

Procedure

- 1 (Optional) If you have a Windows operating system, run `ep-agent.bat remove` to remove the agent service.
- 2 Select the uninstall option that is appropriate to your situation.
 - If you do not intend to reinstall the agent after you have uninstalled it, delete the agent directory.
The default name of the directory is `epops-agent-version`.
 - If you are reinstalling the agent after you have uninstalled it, delete the `/data` directory.
- 3 (Optional) If you do not intend to reinstall the agent after you have uninstalled it, or you do not need to maintain data continuity, delete the `epops-token` platform token file.
Depending on your operating system, the file to delete is one of the following, unless otherwise defined in the properties file.
 - Linux: `/etc/epops/epops-token`
 - Windows: `%PROGRAMDATA%/VMware/EP Ops Agent/epops-token`

Uninstall an Agent that was Installed Using an RPM Package

You can use this procedure to uninstall agents that you installed on virtual machines in your environment using an RPM package.

When you are uninstalling an End Point Operations Management agent, it is good practice to stop the agent running, to reduce unnecessary load on the server.

Procedure

- ◆ On the virtual machine from which you are removing the agent, open a command line and run `rpm -e epops-agent`.

Results

The agent is uninstalled from the virtual machine.

Uninstall an Agent that was Installed Using a Windows Executable

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from a Windows EXE file.

When you are uninstalling an End Point Operations Management agent, it is good practice to stop the agent running, to reduce unnecessary load on the server.

Procedure

- ◆ Double-click `unins000.exe` in the installation destination directory for the agent.

Results

The agent is uninstalled from the virtual machine.

Reinstall an Agent

If you change the IP address, hostname or port number of the vRealize Operations Manager server, you need to uninstall and reinstall your agents.

Prerequisites

To maintain data continuity, you must have retained the `epops-token` platform token file when you uninstalled your agent. See [Uninstall an Agent that was Installed from an Archive](#).

When you reinstall an End Point Operations Management agent on a virtual machine, objects that had previously been detected are no longer monitored. To avoid this situation, do not restart the End Point Operations Management agent until the plug-in synchronization is complete.

Procedure

- ◆ Run the agent install procedure that is relevant to your operating system.
See [Selecting an Agent Installer Package](#).

What to do next

After you reinstall an agent, MSSQL resources might stop receiving data. If this happens, edit the problematic resources and click **OK**.

Install Multiple End Point Operations Management Agents Simultaneously

If you have multiple End Point Operations Management agents to install at one time, you can create a single standardized `agent.properties` file that all the agents can use.

Installing multiple agents entails a number of steps. Perform the steps in the order listed.

Prerequisites

Verify that the following prerequisites are satisfied.

- 1 Set up an installation server.

An installation server is a server that can access the target platforms from which to perform remote installation.

The server must be configured with a user account that has permissions to SSH to each target platform without requiring a password.

- 2 Verify that each target platform on which an End Point Operations Management agent will be installed has the following items.
 - A user account that is identical to that created on the installation server.
 - An identically named installation directory, for example `/home/epomagent`.
 - A trusted keystore, if required.

Procedure

1 [Create a Standard End Point Operations Management Agent Properties File](#)

You can create a single properties file that contains property values that multiple agents use.

2 [Deploy and Start Multiple Agents One-By-One](#)

You can perform remote installations to deploy multiple agents that use a single `agent.properties` file one-by-one.

3 [Deploy and Start Multiple Agents Simultaneously](#)

You can perform remote installations to simultaneously deploy agents that use a single `agent.properties` file.

Create a Standard End Point Operations Management Agent Properties File

You can create a single properties file that contains property values that multiple agents use.

To enable multiple agent deployment, you create an `agent.properties` file that defines the agent properties required for the agent to start up and connect with the vRealize Operations Manager server. If you supply the necessary information in the properties file, each agent locates its setup configuration at startup, rather than prompting you for the location. You can copy the agent properties file to the agent installation directory, or to a location available to the installed agent.

Prerequisites

Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.

Procedure

1 Create an `agent.properties` file in a directory.

You will copy this file later to other machines.

2 Configure the properties as required.

The minimum configuration is the IP address, user name, password, thumb print, and port of the vRealize Operations Manager installation server.

3 Save your configurations.

Results

The first time that the agents are started, they read the `agent.properties` file to identify the server connection information. The agents connect to the server and register themselves.

What to do next

Perform remote agent installations. See [Deploy and Start Multiple Agents One-By-One](#) or [Deploy and Start Multiple Agents Simultaneously](#).

Deploy and Start Multiple Agents One-By-One

You can perform remote installations to deploy multiple agents that use a single `agent.properties` file one-by-one.

Prerequisites

- Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.
- Verify that you configured a standard agent properties file and copied it to the agent installation, or to a location available to the agent installation.

Procedure

- 1 Log in to the installation server user account that you configured with permissions to use SSH to connect to each target platform without requiring a password.
- 2 Use SSH to connect to the remote platform.
- 3 Copy the agent archive to the agent host.
- 4 Unpack the agent archive.
- 5 Copy the `agent.properties` file to the `AgentHome/conf` directory of the unpacked agent archive on the remote platform.
- 6 Start the new agent.

Results

The agent registers with the vRealize Operations Manager server and the agent runs an autodiscovery scan to discover its host platform and supported managed products that are running on the platform.

Deploy and Start Multiple Agents Simultaneously

You can perform remote installations to simultaneously deploy agents that use a single `agent.properties` file.

Prerequisites

- Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.

- Verify that you configured a standard agent properties file and copied it to the agent installation, or to a location available to the agent installation. See [Create a Standard End Point Operations Management Agent Properties File](#).

Procedure

- 1 Create a `hosts.txt` file on your installation server that maps the hostname to the IP address of each platform on which you are installing an agent.
- 2 Open a command-line shell on the installation server.
- 3 Type the following command in the shell, supplying the correct name for the agent package in the `export` command.

```
$ export AGENT=epops-agent-x86-64-linux-1.0.0.tar.gz
$ export PATH_TO_AGENT_INSTALL=</path/to/agent/install>
$ for host in `cat hosts.txt`; do scp $AGENT $host:$PATH_TO_AGENT_INSTALL && ssh $host "cd
$PATH_TO_AGENT_INSTALL; tar xzfp $AGENT &&
./epops-agent-1.0.0/ep-agent.sh start"; done
```

- 4 (Optional) If the target hosts have sequential names, for example `host001`, `host002`, `host003`, and so on, you can skip the `hosts.txt` file and use the `seq` command.

```
$ export AGENT=epops-agent-x86-64-linux-1.0.0.tar.gz
$ for i in `seq 1 9`; do scp $AGENT host$i: && ssh host$i "tar xzfp $AGENT &&
./epops-agent-1.0.0/ep-agent.sh start"; done
```

Results

The agents register with the vRealize Operations Manager server and the agents run an autodiscovery scan to discover their host platform and supported managed products that are running on the platform.

Upgrade the End Point Operations Management Agent

You can upgrade the 6.3 or 6.4 version of an End Point Operations Management agent to a 6.5 version or later, from the vRealize Operations Manager administration interface.

Prerequisites

- Download the End Point Operations Management PAK file.
- Before you install the PAK file, or upgrade your vRealize Operations Manager instance, clone any customized content to preserve it. Customized content can include alert definitions, symptom definitions, recommendations, and views. Then, during the software update, you select the options named **Install the PAK file even if it is already installed** and **Reset out-of-the-box content**.

Procedure

- 1 Log into the vRealize Operations Manager administration interface of your cluster at `https://IP-address/admin`.

- 2 Click **Software Update** in the left panel.
- 3 Click **Install a Software Update** in the main panel.
- 4 From the **Add Software Update** dialog box, click **Browse** to select the PAK file.
- 5 Click **Upload** and follow the steps in the wizard to install your PAK file.
- 6 After Step 4 of the install is complete, you return to the Software Update page of the End Point Operations Management administration interface.
- 7 A message that indicates that the software update completed successfully appears in the main pane.

If any of the agents have not installed successfully, rerun the upgrade steps and ensure that you have selected **Install the PAK file even if it is already installed** in the Add Software Update - Select Software Update page.

What to do next

You can view the log files from the vRealize Operations Manager administration interface > Support page.

Access and View the Log Files

You can access and view the log files to troubleshoot agent upgrade failure. You can verify the status of the agents during and after the upgrade process to find out if the agents have upgraded successfully.

You can view the status of the agents during the upgrade from the `epops-agent-upgrade-status.txt` file. You can view a final report of the number of agents that have successfully upgraded or failed upgrade from the `epops-agent-bundle-upgrade-summary.txt` file.

Procedure

- 1 Log into the vRealize Operations Manager administration interface of your cluster at `https://IP-address/admin`.
- 2 Click **Support** in the left panel.
- 3 Click the **Logs** tab in the right pane and double-click **EPOPS**.
- 4 Double-click the log file to view the contents.

Roles and Privileges in vRealize Operations Manager

vRealize Operations Manager provides several predefined roles to assign privileges to users. You can also create your own roles.

You must have privileges to access specific features in the vRealize Operations Manager user interface. The roles associated with your user account determine the features you can access and the actions you can perform.

Each predefined role includes a set of privileges for users to perform, create, read, update, or delete actions on components such as dashboards, reports, administration, capacity, policies, problems, symptoms, alerts, user account management, and adapters. For information about roles and associated permissions, see [KB 59484](#).

Administrator

Includes privileges to all features, objects, and actions in vRealize Operations Manager.

PowerUser

Users have privileges to perform the actions of the Administrator role except for privileges to user management and cluster management. vRealize Operations Manager maps vCenter Server users to this role.

PowerUserMinusRemediation

Users have privileges to perform the actions of the Administrator role except for privileges to user management, cluster management, and remediation actions.

ContentAdmin

Users can manage all content, including views, reports, dashboards, and custom groups in vRealize Operations Manager.

AgentManager

Users can deploy and configure End Point Operations Management agents.

GeneralUser-1 through GeneralUser-4

These predefined template roles are initially defined as ReadOnly roles. vCenter Server administrators can configure these roles to create combinations of roles to give users multiple types of privileges. Roles are synchronized to vCenter Server once during registration.

ReadOnly

Users have read-only access and can perform read operations, but cannot perform write actions such as create, update, or delete.

Registering Agents on Clusters

You can streamline the process of registering agents on clusters by defining a DNS name for a cluster and configuring that cluster so that the metrics are shared sequentially in a loop.

You only need to register the agent on the DNS, not on the IP address of each individual machine in the cluster. If you do register the agent on each node in the cluster, it affects the scale of your environment.

When you have configured the cluster so that the received metrics are shared in a sequential loop, each time that the agent queries the DNS server for an IP address, the returned address is for one of the virtual machines in the cluster. The next time the agent queries the DNS, it sequentially supplies the IP address of the next virtual machine in the cluster, and so on. The clustered machines are set up in a loop configuration so that each machine receives metrics in turn, ensuring a balanced load.

After you configure the DNS, it is important to maintain it, ensuring that when machines are added or removed from the cluster, their IP address information is updated accordingly.

Manually Create Operating System Objects

The agent discovers some of the objects to monitor. You can manually add other objects, such as files, scripts or processes, and specify the details so that the agent can monitor them.

The **Monitor OS Object** action only appears in the **Actions** menu of an object that can be a parent object.

Procedure

- 1 In the left pane of vRealize Operations Manager, select the agent adapter object that is to be the parent under which you are creating an OS object.

- 2 Select **Actions > Monitor OS Object**.

A list of parent object context-sensitive objects appear in the menu.

- 3 Choose one of the following options.

- Click an object type from the list to open the Monitor OS Object dialog box for that object type.

The three most popularly selected object types appear in the list.

- If the object type that you want to select is not in the list, click **More** to open the Monitor OS Object dialog box. Select the object type from the complete list of objects that are available for selection in the **Object Type** menu.

- 4 Specify a display name for the OS object.
- 5 Enter the appropriate values in the other text boxes.

The options in the menu are filtered according to the OS object type that you select.

Some text boxes might display default values, which you can overwrite if necessary. Note the following information about default values.

Option	Value
Process	<p>Supply the PTQL query in the form: <code>Class.Attribute.operator=value</code>. For example, <code>Pid.PidFile.eq=/var/run/sshd.pid</code>.</p> <p>Where:</p> <ul style="list-style-type: none"> ■ <code>Class</code> is the name of the Sigar class without the Proc prefix. ■ <code>Attribute</code> is an attribute of the given Class, index into an array or key in a Map class. ■ <code>operator</code> is one of the following (for String values): <ul style="list-style-type: none"> ■ <code>eq</code> Equal to value ■ <code>ne</code> Not Equal to value ■ <code>ew</code> Ends with value ■ <code>sw</code> Starts with value ■ <code>ct</code> Contains value (substring) ■ <code>re</code> Regular expression value matches <p>Delimit queries with a comma.</p>
Windows Service	<p>Monitor an application that runs as a service under Windows.</p> <p>To configure it, you supply its Service Name in Windows.</p> <p>To determine the Service Name:</p> <ol style="list-style-type: none"> 1 Select Run from the Windows Start menu. 2 Type <code>services.msc</code> in the run dialog box and click OK. 3 In the list of services displayed, right-click the service to monitor and choose Properties. 4 Locate the Service Name on the General tab.
Script	<p>Configure vRealize Operations Manager to periodically run a script that collects a system or application metric.</p>

6 Click **OK**.

You cannot click **OK** until you enter values for all the mandatory text boxes.

Results

The OS object appears under its parent object and monitoring begins.

Caution If you enter invalid details when you create an OS object, the object is created but the agent cannot discover it, and metrics are not collected.

Managing Objects with Missing Configuration Parameters

Sometimes when an object is discovered by vRealize Operations Manager for the first time, the absence of values for some mandatory configuration parameters is detected. You can edit the object's parameters to supply the missing values.

If you select **Custom Groups > Objects with Missing Configuration (EP Ops)** in the Environment Overview view of vRealize Operations Manager, you can see the list of all objects that have missing mandatory configuration parameters. In addition, objects with such missing parameters return an error in the Collection Status data.

If you select an object in the vRealize Operations Manager user interface that has missing configuration parameters, the red Missing Configuration State icon appears on the menu bar. When you point to the icon, details about the specific issue appear.

You can add the missing parameter values through the **Action > Edit Object** menu.

Mapping Virtual Machines to Operating Systems

You can map your virtual machines to an operating system to provide additional information to assist you to determine the root cause of why an alert was triggered for a virtual machine.

vRealize Operations Manager monitors your ESXi hosts and the virtual machines located on them. When you deploy an End Point Operations Management agent, it discovers the virtual machines and the objects that are running on them. By correlating the virtual machines discovered by the End Point Operations Management agent with the operating systems monitored by vRealize Operations Manager you have more details to determine the exact cause of an alert being triggered.

Verify that you have the vCenter Adapter configured with the vCenter Server that manages the virtual machines. You also need to ensure that you have VMware Tools that are compatible with the vCenter Server installed on each of the virtual machines.

User Scenario

vRealize Operations Manager is running but you have not yet deployed the End Point Operations Management agent in your environment. You configured vRealize Operations Manager to send you alerts when CPU problems occur. You see an alert on your dashboard because insufficient CPU capacity is available on one of your virtual machines that is running a Linux operating system. You deploy another two virtual CPUs but the alert remains. You struggle to determine what is causing the problem.

In the same situation, if you deployed the End Point Operations Management agent, you can see the objects on your virtual machines, and determine that an application-type object is using all available CPU capacity. When you add more CPU capacity, it also uses that. You disable the object and your CPU availability is no longer a problem.

Viewing Objects on Virtual Machines

After you deploy an End Point Operations Management agent on a virtual machine, the machine is mapped to the operating system and you can see the objects on that machine.

All the actions and the views that are available to other objects in your vRealize Operations Manager environment are also available for newly discovered server, service, and application objects, and for the deployed agent.

You can see the objects on a virtual machine in the inventory when you select the machine by clicking **Environment** from the menu, and then from the left pane click **vSphere Environment > vSphere Hosts and Clusters**. You can see the objects and the deployed agent under the operating system.

When you select an object, the center pane of the user interface displays data relevant to that objects.

Customizing How End Point Operations Management Monitors Operating Systems

End Point Operations Management gathers operating system metrics through agent-based collections. In addition to the features available after initial configuration of End Point Operations Management, you can enable remote monitoring, enable or disable plug-ins for additional monitoring, and customize End Point Operations Management logging.

Configuring Remote Monitoring

With remote monitoring you can monitor the state of an object from a remote location by configuring a remote check.

You can configure remote monitoring using HTTP, ICMP TCP methods.

When you configure a remote HTTP, ICMP or TCP check, it is created as a child object of the tested object that you are monitoring and of the monitoring agent.

If the object that you select to remotely monitor does not already have an alert configured, one is created automatically in the format *Remote check type failed on a object type*. If the object has an existing alert, that is used.

Configure Remote Monitoring of an Object

Use this procedure to configure remote monitoring of an object.

Configuration options are defined in [HTTP Configuration Options](#), [ICMP Configuration Options](#) and [TCP Configuration Options](#). You might need to refer to this information when you are completing this procedure.

Procedure

- 1 In the vRealize Operations Manager user interface, select the remote object to monitor.
- 2 On the details page for the object, select **Monitor this Object Remotely** from the **Actions** menu.
- 3 In the Monitor Remote Object dialog, select the End Point Operations Management agent that will remotely monitor the object from the **Monitored From** menu.
- 4 Select the method with which the remote object will be monitored from the **Check Method** menu.

The relevant parameters for the selected object type appear.

- 5 Enter values for all of the configuration options and click **OK**.

HTTP Configuration Options

Here are the options in the configuration schema for the HTTP resource.

For the HTTP resource, the netsservices plug-in descriptor default values are:

- port: 80
- sslport: 443

HTTP Configuration Options

Table 1-4. ssl Option

Option Information	Value
Description	Use ssl
Default	false
Optional	true
Type	boolean
Notes	N/A
Parent Schema	ssl

Table 1-5. hostname Option

Option Information	Value
Description	Hostname
Default	localhost
Optional	false
Type	N/A
Notes	The hostname of system that hosts the service to monitor. For example: mysite.com
Parent Schema	sockaddr

Table 1-6. port Option

Option Information	Value
Description	Port
Default	A default value for port is set for each type of network service by properties in the netsservices plug-in descriptor.
Optional	false
Type	N/A
Notes	The port on which the service listens.
Parent Schema	sockaddr

Table 1-7. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10

Table 1-7. sotimeout Option (continued)

Option Information	Value
Optional	true
Type	int
Notes	The maximum length of time the agent waits for a response to a request to the remote service.
Parent Schema	sockaddr

Table 1-8. path Option

Option Information	Value
Description	Path
Default	/
Optional	false
Type	N/A
Notes	Enter a value to monitor a specific page or file on the site. for example: /Support.html.
Parent Schema	url

Table 1-9. method Option

Option Information	Value
Description	Request Method
Default	HEAD
Optional	false
Type	enum
Notes	Method for checking availability. Permitted values: HEAD, GET HEAD results in less network traffic. Use GET to return the body of the request response to specify a pattern to match in the response.
Parent Schema	http

Table 1-10. hostheader Option

Option Information	Value
Description	Host Header
Default	none
Optional	true
Type	N/A

Table 1-10. hostheader Option (continued)

Option Information	Value
Notes	Use this option to set a Host HTTP header in the request. This is useful if you use name-based virtual hosting. Specify the host name of the Vhost's host, for example, blog.mypost.com.
Parent Schema	http

Table 1-11. follow Option

Option Information	Value
Description	Follow Redirects
Default	enabled
Optional	true
Type	boolean
Notes	Enable if the HTTP request that is generated will be redirected. This is important, because an HTTP server returns a different code for a redirect and vRealize Operations Manager determines that the HTTP service check is unavailable if it is a redirect, unless this redirect configuration is set.
Parent Schema	http

Table 1-12. pattern Option

Option Information	Value
Description	Response Match (substring or regex)
Default	none
Optional	true
Type	N/A
Notes	Specify a pattern or substring for vRealize Operations Manager to attempt to match against the content in the HTTP response. This enables you to check that in addition to being available, the resource is serving the content you expect.
Parent Schema	http

Table 1-13. proxy Option

Option Information	Value
Description	Proxy Connection
Default	none
Optional	true
Type	N/A

Table 1-13. proxy Option (continued)

Option Information	Value
Notes	If the connection to the HTTP service goes through a proxy server, supply the hostname and port for the proxy server. For example, proxy.myco.com:3128.
Parent Schema	http

Table 1-14. requestparams Option

Option Information	Value
Description	Request arguments. For example, arg0=val0, arg1=val1, and so on.
Default	N/A
Optional	true
Type	string
Notes	Request parameters added to the URL to be tested.
Parent Schema	http

Table 1-15. Credential Option

Option Information	Value
Description	Username
Default	N/A
Optional	true
Type	N/A
Notes	Supply the user name if the target site is password-protected.
Parent Schema	credentials

ICMP Configuration Options

Here are the options in the configuration schema for the ICMP resource.

ICMP configuration is not supported in Windows environments. When attempting to run an ICMP check for remote monitoring from an Agent running on a Windows platform, no data is returned.

Table 1-16. hostname Option

Option Information	Value
Description	Hostname
Default	localhost
Optional	N/A
Type	N/A

Table 1-16. hostname Option (continued)

Option Information	Value
Notes	The hostname of system that hosts the object to monitor. For example: mysite.com
Parent Schema	netservices plug-in descriptor

Table 1-17. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A
Type	int
Notes	The maximum time period the agent waits for a response to a request to the remote service.
Parent Schema	netservices plug-in descriptor

TCP Configuration Options

Here are the options in the configuration schema to enable TCP checking.

Table 1-18. port Option

Option Information	Value
Description	Port
Default	A default value for port is set for each type of network service by properties in the netservices plug-in descriptor.
Optional	false
Type	N/A
Notes	The port on which the service listens.
Parent Schema	sockaddr

Table 1-19. hostname Option

Option Information	Value
Description	Hostname
Default	localhost
Optional	N/A
Type	N/A
Notes	The hostname of system that hosts the object to monitor. For example: mysite.com
Parent Schema	netservices plug-in descriptor

Make sure that you use the IP address of the machine on which the remote check is to run, not the host name.

Table 1-20. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A
Type	int
Notes	The maximum amount of time the agent waits for a response to a request to the remote service.
Parent Schema	netservices plug-in descriptor

Agent Management

You can add, edit, and delete End Point Operations Management agents and enable or disable the End Point Operations Management plug-ins from the tabs in the Agent Management page.

Where You Find the Agent Management Page

In the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**.

Agents Tab

You can view the End Point Operations Management agents that are installed and deployed in your environment.

Where You Find the Agents Tab

In the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**.

How the Agents Tab Works

You can view all the agents that are installed, the virtual machines on which they are installed, their operating system and the agent bundle version. You can also view the collection details of each agent. You can filter the list of agents based on the name of the agent. You add a filter from the upper-right corner of the toolbar. You can sort the Agent Token, Agent Name, Collection State, and Collection Status columns by clicking the column name.

Plug-ins Tab

End Point Operations Management agents include plug-ins that determine which objects to monitor, how they should be monitored, which metrics to collect, and so on. Some plug-ins are included in the default End Point Operations Management agent installation, and other plug-ins might be added as part of any management pack solution that you install to extend the vRealize Operations Manager monitoring process.

You can use the **Plug-ins** tab from the Agents Management page to disable or enable the agent plug-ins that are deployed in your environment as part of a solution installation. For example, you might want to temporarily disable a plug-in so that you can analyze the implication of that plug-in on a monitored virtual machine. To access the **Plug-ins** tab, in the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**. You can sort all the columns in the tab by clicking the column name.

All the default plug-ins and the plug-ins that are deployed when you installed one or more solutions are listed alphabetically on the tab.

You must have Manage Plug-ins permissions to enable and disable plug-ins.

When you disable a plug-in, it is removed from all the agents on which it has existed, and the agent no longer collects the metrics and other data related to that plug-in. The plug-in is marked as disabled on the vRealize Operations Manager server.

You cannot disable the default plug-ins that are installed during the vRealize Operations Manager installation.

You use the action menu that appears when you click the gear wheel icon to disable or enable plug-ins.

Before you deploy a new version of a plug-in, you must implement a shutdown method. If you do not implement a shutdown method, the existing plug-in version does not shut down so that a new instance is created and allocated resources such as static threads are not released. Implement a shutdown method for these plug-ins.

- Plug-ins that use third-party libraries
- Plug-ins that use native libraries
- Plug-ins that use connection pools
- Plug-ins that might lock files, which cause issues on Windows operating systems

It is good practice that plug-ins do not use threads, third-party libraries, or static collection.

Configuring Plug-in Loading

At startup, an End Point Operations Management agent loads all the plug-ins in the AgentHome/bundles/agent-x.y.z-nnnn/pdk/plugins directory. You can configure properties in the agent.properties file to reduce an agent's memory footprint by configuring it to load only the plug-ins that you use.

Plug-ins are deployed to all agents when a solution is installed. You might want to use the properties described here in a situation in which you need to remove one or more plug-ins from a specific machine. You can either specify a list of plug-ins to exclude, or configure a list of plug-ins to load.

plugins.exclude

Use this property to specify the plug-ins that the End Point Operations Management agent must not load at startup.

You supply a comma-separated list of plugins to exclude. For example,
`plugins.exclude=jboss,apache,mysql.`

plugins.include

Use this property to specify the plug-ins that the End Point Operations Management agent must load at startup.

You supply a comma-separated list of plugins to include. For example,
`plugins.include=weblogic,apache.`

Understanding the Unsynchronized Agents Group

An unsynchronized agent is an agent that is not synchronized with the vRealize Operations Manager server in terms of its plug-ins. The agent might be missing plug-ins that are registered on the server, include plug-ins that are not registered on the server, or include plug-ins that have a different version to that registered on the server.

Each agent must be synchronized with the vRealize Operations Manager server. During the time that an agent is not synchronized with the server, it appears in the Unsynchronized Agents list. The list is located in the vRealize Operations Manager user interface on the **Groups** tab in the Environment view.

The first time an agent is started, a status message is sent to the server. The server compares the status sent by the agent with that on the server. The server sends commands to the agent to synchronize, download or delete plug-ins, as required by the differences that it detects.

When a plug-in is deployed, disabled, or enabled as part of a management pack solution update, the vRealize Operations Manager server detects that change and sends a new command to the agents so that synchronization occurs.

Commonly, multiple agents are affected at the same time when a plug-in is deployed, disabled or enabled. All agents have an equal need to be updated so, to avoid overloading the server and creating performance issues that might occur if many agents were all synchronized at the same time, synchronization is performed in batches and is staggered in one-minute periods. You will notice that the list of unsynchronized agents decrements over time.

Configuring Agent Logging

You can configure the name, location, and logging level for End Point Operations Management agent logs. You can also redirect system messages to the agent log, and configure the debug log level for an agent subsystem.

Agent Log Files

The End Point Operations Management agent log files are stored in the `AgentHome/log` directory.

Agent log files include the following:

agent.log

agent.operations.log

This log is applicable to Windows-based agents only.

This is an audit log that records the commands that were run on the agent, together with the parameters that the agent used to action them.

wrapper.log

The Java service wrapper-based agent launcher writes messages to the `wrapper.log` file. For a non-JRE agent, this file is located in `agentHome/wrapper/sbin`.

In the event that the value was changed ifr the `agent.logDir` property, the file is also located in `agentHome/wrapper/sbin`.

Configuring the Agent Log Name or Location

Use these properties to change the name or location of the agent log file.

agent.logDir

You can add this property to the `agent.properties` file to specify the directory where the End Point Operations Management agent will write its log file. If you do not specify a fully qualified path, `agent.logDir` is evaluated relative to the agent installation directory.

This property does not exist in the `agent.properties` file unless you explicitly add it. The default behavior is equivalent to the `agent.logDir=log` setting, resulting in the agent log file being written to the `AgentHome/log` directory.

To change the location for the agent log file, add `agent.logDir` to the `agent.properties` file and enter a path relative to the agent installation directory, or a fully qualified path.

The name of the agent log file is configured with the `agent.logFile` property.

agent.logFile

This property specifies the path and name of the agent log file.

In the `agent.properties` file, the default setting for the `agent.LogFile` property is made up of a variable and a string, `agent.logFile=${agent.logDir}\agent.logDir`.

- *agent.logDir* is a variable that supplies the value of an identically named agent property. By default, the value of *agent.logDir* is `log`, interpreted relative to the agent installation directory.
- `agent.log` is the name for the agent log file.

By default, the agent log file is named `agent.log` and is written to the `AgentHome/log` directory.

To configure the agent to log to a different directory, you must explicitly add the `agent.logDir` property to the `agent.properties` file.

Configuring the Agent Logging Level

Use this property to control the severity level of messages that the End Point Operations Management agent writes to the agent log file.

agent.logLevel

This property specifies the level of detail of the messages that the End Point Operations Management agent writes to the log file.

Setting the `agent.logLevel` property value to `DEBUG` level is not advised. This level of logging across all subsystems imposes overhead, and can also cause the log file to roll over so frequently that log messages of interest are lost. It is preferable to configure debug level logging only at the subsystem level.

The changes that you make to this property become effective approximately five minutes after you save the properties file. It is not necessary to restart the agent to initiate the change.

Redirecting System Messages to the Agent Log

You can use these properties to redirect system-generated messages to the End Point Operations Management agent log file.

agent.logLevel.SystemErr

This property redirects `System.err` to `agent.log`. Commenting out this setting causes `System.err` to be directed to `agent.log.startup`.

The default value is `ERROR`.

agent.logLevel.SystemOut

This property redirects `System.out` to `agent.log`. Commenting out this setting causes `System.out` to be directed to `agent.log.startup`.

The default value is `INFO`.

Configuring the Debug Level for an Agent Subsystem

For troubleshooting purposes, you can increase the logging level for an individual agent subsystem.

To increase the logging level for an individual agent subsystem, uncomment the appropriate line in the section of the `agent.properties` file that is labeled `Agent Subsystems: Uncomment individual subsystems to see debug messages`.

Agent log4j Properties

This is the `log4j` properties in the `agent.properties` file.

```
log4j.rootLogger=${agent.logLevel}, R

log4j.appender.R.File=${agent.logFile}
log4j.appender.R.MaxBackupIndex=1
log4j.appender.R.MaxFileSize=5000KB
log4j.appender.R.layout.ConversionPattern=%d{dd-MM-yyyy HH:mm:ss,SSS z} %-5p [%t] [%c{1}@%L] %m%n
log4j.appender.R.layout=org.apache.log4j.PatternLayout
log4j.appender.R=org.apache.log4j.RollingFileAppender

##
## Disable overly verbose logging
```



```

##
log4j.logger.org.apache.http=ERROR
log4j.logger.org.springframework.web.client.RestTemplate=ERROR
log4j.logger.org.hyperic.hq.measurement.agent.server.SenderThread=INFO
log4j.logger.org.hyperic.hq.agent.server.AgentDListProvider=INFO
log4j.logger.org.hyperic.hq.agent.server.MeasurementSchedule=INFO
log4j.logger.org.hyperic.util.units=INFO
log4j.logger.org.hyperic.hq.product.pluginxml=INFO

# Only log errors from naming context
log4j.category.org.jnp.interfaces.NamingContext=ERROR
log4j.category.org.apache.axis=ERROR

#Agent Subsystems: Uncomment individual subsystems to see debug messages.
#-----
#log4j.logger.org.hyperic.hq.autoinventory=DEBUG
#log4j.logger.org.hyperic.hq.livedata=DEBUG
#log4j.logger.org.hyperic.hq.measurement=DEBUG
#log4j.logger.org.hyperic.hq.control=DEBUG

#Agent Plugin Implementations
#log4j.logger.org.hyperic.hq.product=DEBUG

#Server Communication
#log4j.logger.org.hyperic.hq.bizapp.client.AgentCallbackClient=DEBUG

#Server Realtime commands dispatcher
#log4j.logger.org.hyperic.hq.agent.server.CommandDispatcher=DEBUG

#Agent Configuration parser
#log4j.logger.org.hyperic.hq.agent.AgentConfig=DEBUG

#Agent plugins loader
#log4j.logger.org.hyperic.util.PluginLoader=DEBUG

#Agent Metrics Scheduler (Scheduling tasks definitions & executions)
#log4j.logger.org.hyperic.hq.agent.server.session.AgentSynchronizer.SchedulerThread=DEBUG

#Agent Plugin Managers
#log4j.logger.org.hyperic.hq.product.MeasurementPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.AutoinventoryPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ConfigTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LogTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LiveDataPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ControlPluginManager=DEBUG

```

VMware vRealize Application Management Pack

The VMware vRealize Application Management Pack enables application monitoring from Wavefront.

Do not add, edit, or modify operations. For information about adding an application proxy, see [Add and Configure an Application Proxy](#).

View the Configuration Details

You can view configuration details of the VMware vRealize Application Management Pack.

To access and view the configuration details, complete the following steps:

- 1 In the menu, select **Administration**, and then from the left pane, select **Solutions**.
- 2 From the Solutions pane in the right pane, select VMware vRealize Application Management Pack.
- 3 Click the **Configure** icon.

Table 1-21. Configuration Details

Options	Description
Instance Name	Displays the vCenter servers that have been mapped with the VMware Application Proxy.
Display Name	Displays the IP address of the VMware Application Proxy and the vCenter Server.
UCP Host	Displays the IP address of the VMware Application Proxy you have configured.
Mapped vCenter(s)	Displays the IP address of the vCenter Server you mapped to the VMware Application Proxy.
Credentials	<p>Displays the name of the credential, which is the IP address of the VMware Application Proxy.</p> <p>To add credentials, click the plus sign.</p> <ul style="list-style-type: none"> ■ Credential Name: The name by which you are identifying and managing the configured credentials. ■ Application Proxy Username: The user account details used in VMware Application Proxy. ■ Application Proxy Password: Password of the user account in VMware Application Proxy.
Collectors/Groups	Select the collector that is used to manage the adapter processes.

Log Insight

When vRealize Operations Manager is integrated with Log Insight, you can view the Log Insight page, the Troubleshoot with Logs dashboard, and the Logs tab. You can collect and analyze log feeds. You can filter and search for log messages. You can also dynamically extract fields from log messages based on customized queries.

Log Insight Page

When vRealize Operations Manager is integrated with vRealize Log Insight, you can search and filter log events. From the Interactive Analytics tab in the Log Insight page, you can create queries to extract events based on timestamp, text, source, and fields in log events. vRealize Log Insight presents charts of the query results.

To access the Log Insight page from vRealize Operations Manager, you must either:

- Configure the vRealize Log Insight adapter from the vRealize Operations Manager interface, or
- Configure vRealize Operations Manager in vRealize Log Insight.

For more information about configuring, see [Configuring vRealize Log Insight with vRealize Operations Manager](#).

For information about vRealize Log Insight interactive analytics, see the [vRealize Log Insight documentation](#).

Logs Tab

When vRealize Operations Manager is integrated with vRealize Log Insight, you can view the logs for a selected object from the Logs tab. You can troubleshoot a problem in your environment by correlating the information in the logs with the metrics. You can then most likely determine the root cause of the problem.

How the Logs Tab Works

By default, the Logs tab displays different event types for the last hour. For vSphere objects, the logs are filtered to show the event types for the specific object you select. For more information on the different filtering and querying capabilities, see the [vRealize Log Insight documentation](#).

Where You Find the Logs Tab

In the menu, select **Environment** and then from the left pane select an inventory object. Click the **Logs** tab. To view the Logs tab, you have to configure vRealize Operations Manager in vRealize Log Insight. For more information, see [Configuring vRealize Log Insight with vRealize Operations Manager](#).

After integrating vRealize Operations Manager with vRealize Log Insight, refresh the browser to see the Logs tab.

Configuring vRealize Log Insight with vRealize Operations Manager

To use the Log Insight page, the Troubleshoot with Logs dashboard, and Logs tab in vRealize Operations Manager, you must configure vRealize Log Insight with vRealize Operations Manager.

Configuring the vRealize Log Insight Adapter in vRealize Operations Manager

To access the Log Insight page and the Troubleshoot with Logs dashboard from vRealize Operations Manager, you must configure the vRealize Log Insight adapter in vRealize Operations Manager.

vRealize Operations Manager accesses the first instance of the vRealize Log Insight adapter that is configured.

Prerequisites

- Verify that vRealize Log Insight and vRealize Operations Manager are installed.

- Verify that you know the IP address, user name, and password of the vRealize Log Insight instance you have installed.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Solutions**.
- 2 From the Solutions page, click VMware vRealize Log Insight.
- 3 Click the **Configure** icon. You see the Manage Solution-VMware vRealize Log Insight dialog box.
- 4 In the Manage Solutions dialog box perform the following steps:
 - Enter a name in the **Display Name** text box.
 - Enter the IP address in the **Log Insight server** text box of the vRealize Log Insight you have installed and want to integrate with.
 - Click **Test Connection** to verify that the connection is successful.
 - Click **Save Settings**.
 - Click **Close**.
- 5 From the vRealize Operations Manager Home page, click **Troubleshoot > Using Logs** from the left pane. If you see a statement at the bottom of the page, click the link and accept the certificate exception in vRealize Log Insight or contact your IT support for more information.
- 6 From the vRealize Operations Manager Home page, click **Troubleshoot > Using Logs** from the left pane and enter the user name and password of the vRealize Log Insight instance you have installed.

Configuring vRealize Operations Manager in vRealize Log Insight

You configure vRealize Operations Manager in vRealize Log Insight in the following scenarios:

- To access the Logs tab in vRealize Operations Manager.
- To access the Troubleshoot with Logs dashboard and the Log Insight page from vRealize Operations Manager.

Prerequisites

- Verify that vRealize Log Insight and vRealize Operations Manager are installed.
- Verify that you know the IP address, hostname, and password of the vRealize Operations Manager instance you want to integrate with.

Procedure

- 1 From the Administration page of vRealize Log Insight, click the **vRealize Operations** icon from the left pane. You see the vRealize Operations Integration pane.
- 2 In the **Hostname** and **Username** text boxes, enter the IP address and hostname of the vRealize Operations Manager instance you want to integrate with.

- 3 In the **Password** text box, select **Update Password** and enter the password of the vRealize Operations Manager instance you want to integrate with.
- 4 Select the **Enable launch in context** option.
- 5 Click **Test Connection** to verify that the connection is successful.
- 6 Click **Save**.

You can now view the log details for an object in vRealize Operations Manager.

Log Forwarding

For troubleshooting in the product UI, you can send the logs to an external log server or a vRealize Log Insight server.

If you have configured log forwarding from **Administration > Support > Logs** in earlier versions of vRealize Operations Manager, VMware recommends that you reconfigure in this version of vRealize Operations Manager.

Where You Find the Log Forwarding Page

In the menu, select **Administration** and then from the left pane select **Management > Log Forwarding**.

Table 1-22. Log Forwarding Page Options

Options	Description															
Output logs to the external log server	Forwards the logs to an external log server.															
Forwarded Logs	You can select the set of logs you want to forward to the external log server or the vRealize Log Insight server.															
Log Insight Servers	You can select an available vRealize Log Insight server IP. If there is no available vRealize Log Insight server IP, select Other from the drop-down menu and manually enter the configuration details.															
Host	IP address of the external log server where logs have to be forwarded.															
Port	<div>The default port value depends on whether or not SSL has been set up for each protocol. The following are the possible default port values:</div> <table><tr><th>Protocol</th><th>SSL</th><th>Default Port</th></tr><tr><td>cfapi</td><td>No</td><td>9000</td></tr><tr><td>cfapi</td><td>Yes</td><td>9543</td></tr><tr><td>syslog</td><td>No</td><td>514</td></tr><tr><td>syslog</td><td>Yes</td><td>6514</td></tr></table>	Protocol	SSL	Default Port	cfapi	No	9000	cfapi	Yes	9543	syslog	No	514	syslog	Yes	6514
Protocol	SSL	Default Port														
cfapi	No	9000														
cfapi	Yes	9543														
syslog	No	514														
syslog	Yes	6514														
Use SSL	Allows the vRealize Log Insight agent to send data securely.															

Table 1-22. Log Forwarding Page Options (continued)

Options	Description
Certificate Path	You can enter the path to the trusted root certificates bundle file. If you do not enter a certificate path, the vRealize Log Insight Windows agent uses system root certificates and the vRealize Log Insight Linux agent attempts to load trusted certificates from <code>/etc/pki/tls/certs/ca-bundle.crt</code> or <code>/etc/ssl/certs/ca-certificates.crt</code> .
Protocol	You can select either <code>cfapi</code> or <code>syslog</code> from the drop-down menu to send event logging messages.
Cluster Name	Displays the name of the cluster. You can edit this field.

Modifying Existing Log Types

If you manually modified the existing entries or logs sections and then modify the log forwarding settings from vRealize Operations Manager, you lose the changes that you made.

The following server entries are overwritten by the vRealize Operations Manager log forwarding settings.

```
port
proto
hostname
ssl
reconnect
ssl_ca_path
```

The following `[common | global]` tags are being added or overwritten by the vRealize Operations Manager log forwarding settings.

```
vmw_vr_ops_appname
vmw_vr_ops_clustername
vmw_vr_ops_clusterrole
vmw_vr_ops_hostname
vmw_vr_ops_nodename
```

Note Cluster role changes do not change the value of the `vmw_vr_ops_clusterrole` tag. You can either manually modify or ignore it.

Business Management

SDDC costing is out-of-the box with vRealize Operations Manager. There is no integration required with vRealize Business for Cloud. You can display infrastructure performance and cost information in the Business Management page.

To display infrastructure performance and cost information, you must configure the vRealize Business for Cloud adapter. For information about configuring this adapter, refer to [Configure the vRealize Business for Cloud Adapter](#).

After you have configured the adapter, you can click the link at the bottom of the Business Management page to log in to vRealize Business for Cloud, and accept the certificate exception.

You must accept the certificate exception each time you log in to vRealize Business for Cloud to see data in the Business Management page.

Configure the vRealize Business for Cloud Adapter

Integrate VMware vRealize Business for Cloud with vRealize Operations Manager to view your infrastructure performance, cost information, and also troubleshooting tips.

You can connect vRealize Operations Manager to a single instance of vRealize Business for Cloud.

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Solutions**.
- 2 Select **VMware vRealize Business for Cloud**, and click the **Configure** icon.
- 3 Enter a name for the adapter instance.
- 4 In the **vRealize Business for Cloud Server** text box, enter the IP address of the vRealize Business for Cloud server to which you want to connect.
- 5 Click **Test Connection** to verify that the connection is successful.
- 6 Click **Advanced Settings**, and in the **Collectors/Groups** text box, select the vRealize Operations Manager collector used to manage the adapter process.

If you have one adapter instance, select **Default collector group**. If you have multiple collectors in your environment, to distribute the workload and optimize performance, select the collector to manage the adapter processes for this instance.

- 7 Click **Save Settings** to finish configuration of the adapter, and click **Close**.

vRealize Business for Cloud adapter is available and is used only as a pre-configuration for SDDC health MP.

Cost Settings for Financial Accounting Model

You can configure Server Hardware cost driver and resource utilization parameters to calculate the accurate cost and improve the efficiency of your environment.

Cost Drivers analyzes the resources and the performance of your virtual environment. Based on the values you define, Cost Drivers can identify reclamation opportunities and can provide recommendations to reduce wastage of resources and cost.

Configuring Depreciation Preferences

To compute the amortized cost of the Server Hardware cost driver, you can configure the depreciation method and the depreciation period. Cost Drivers supports two yearly depreciation methods and you can set the depreciation period from two to seven years.

Note Cost Drivers calculates the yearly depreciation values and then divides the value by 12 to arrive at the monthly depreciation.

Method	Calculation
Straight line	Yearly straight line depreciation = [(original cost – accumulated depreciation) / number of remaining depreciation years]
Max of Double or Straight	Yearly max of Double or Straight = Maximum (yearly depreciation of double declining balance method, yearly depreciation of straight line method) Yearly depreciation of double declining method= [(original cost – accumulated depreciation) * depreciation rate]. Depreciation rate = 2 / number of depreciation years.
Note Double declining depreciation for the last year = original cost – accumulated depreciation	

Example: Example for Straight Line Depreciation Method

Year	Original Cost	Accumulated Depreciation	Straight Line Depreciation Cost
Year 1	10000	0	$[(10000-0)/5] = 2000$
Year 2	10000	2000	$[(10000-2000)/4] = 2000$
Year 3	10000	4000	$[(10000-2000)/3] = 2000$
Year 4	10000	6000	$[(10000-2000)/2] = 2000$
Year 5	10000	8000	$[(10000-2000)/1] = 2000$

Example: Example for Max of Double and Straight Line Depreciation Method

Year	Original Cost	Depreciation Rate	Accumulated Depreciation	Straight Line Depreciation Cost
Year 1	10000	0.4	0	$\text{Maximum}([(10000-0)*0.4], [(10000-0)/5])$ $= \text{Maximum}(4000, 2000) = 4000$ which is 333.33 per month.
Year 2	10000	0.4	4000	$\text{Maximum}([(10000-4000)*0.4], [(10000-4000)/4])$ $= \text{Maximum}(2400, 1500) = 2400$ which is 200 per month.

Year	Original Cost	Depreciation Rate	Accumulated Depreciation	Straight Line Depreciation Cost
Year 3	10000	0.4	6400	$\text{Maximum}([(10000-6400)*0.4], [(10000-6400)/3])$ $= \text{Maximum} (1440, 1200) = 1440$ <p>which is 120 per month.</p>
Year 4	10000	0.4	7840	$\text{Maximum}([(10000-7840)*0.4], [(10000-7840)/2])$ $= \text{Maximum} (864, 1080) = 1080$ <p>which is 90 per month.</p>
Year 5	10000	0.4	8920	$\text{Maximum}([(10000-8920)*0.4], [(10000-8920)/1])$ $= \text{Maximum} (432, 1080) = 1080$ <p>which is 90 per month.</p>

Overview of Cost Drivers

Cost Drivers are the aspect that contributes to the expense of your business operations. Cost drivers provide a link between a pool of costs. To provide a granular cost visibility and to track your expenses of virtual machines accurately in a private cloud, vRealize Operations Manager has identified eight key cost drivers. You can see the total projected expense on your private cloud accounts for the current month and the trend of cost over time.

According to the industry standard, vRealize Operations Manager maintains a reference cost for these cost drivers. This reference cost helps you for calculating the cost of your setup, but might not be accurate. For example, you might have received some special discounts during a bulk purchase or you might have an ELA with VMware that may not match the socket-based pricing available in the reference database. To get accurate values, you can modify the reference cost of cost drivers in vRealize Operations Manager, which overrides the values in reference database. Based on your inputs, vRealize Operations Manager recalculates the total amount of private cloud expenses. After you add a private cloud into vRealize Operations Manager, vRealize Operations Manager automatically discovers one or more vCenter Servers that are part of your Private Cloud. In addition, it also retrieves the inventory details from each vCenter Server. The details include:

- Associated clusters: Count and names
- ESXi hosts: Count, model, configuration, and so on.
- Data stores: Count, storage, type, capacity
- VMs: Count, OS type, tags, configuration, utilization

Based on these configuration and utilizations of inventory, and the available reference cost, vRealize Operations Manager calculates the estimated monthly cost of each cost driver. The total cost of your private cloud is the sum of all these cost driver expenses.

You can modify the expense of your data center. These costs can be in terms of percentage value or unit rate, and might not always be in terms of the overall cost. Based on your inputs, the final amount of expense is calculated. If you do not provide inputs regarding expenses, the default values are taken from the reference database.

You can see the projected cost of private cloud for the current month and the trend of total cost over time. For all the expenses, cost drivers in vRealize Operations Manager display the monthly trend of the cost variations, the actual expense, and a chart that represents the actual expense and the reference cost of the expense.

Note If the vCenter Server was added from more than six months, the trend displays the total cost for the last six months only. Otherwise, the trend displays the total cost from the month the vCenter Server was added into vRealize Operations Manager.

Table 1-23. Expense Types

Cost Drivers	Description
Server Hardware	The Server Hardware cost driver tracks all the expenses for purchasing of hardware servers that are part of vCenter Servers. You see the server cost based on CPU age and server cost details.
Storage	You can calculate the storage cost at the level of a datastore based on tag category information collected from vCenter Server. You see the storage total distribution based on category and the uncategorized cost details.
License	You see the licenses cost distribution for the operating systems cost and VMware license of your cloud environment. Note For Non-ESX physical servers, VMware license is not applicable.
Maintenance	You see the maintenance cost distribution for the server hardware and operating system maintenance. You can track your total expense with hardware and operating system vendors.
Labor	You see the labor cost distribution for the servers, virtual infrastructure, and operating systems. You can view the total administrative cost for managing physical servers, operating systems and virtual machines. You can track all expenses spent on human resources to manage the datacenters. Note <ul style="list-style-type: none"> ■ Labor cost includes expenses on backup appliance virtual machine (VDP virtual appliance). ■ For physical servers, operating system labor cost and servers labor costs are applicable, virtual infrastructure cost is not considered.
Network	You see the networks costs by NIC type. You can track a network expense based on different types of NICs attached to the ESX server. You can view the total cost of physical network infrastructure that includes internet bandwidth, and is estimated by count and type of network ports on the ESXi Servers. Note For physical servers, the network details are not captured. So, the network cost is considered as zero.
Facilities	You see the cost distribution for the facilities such as real estate costs, such as rent or cost of data center buildings, power, cooling, racks, and associated facility management labor cost. You can point to the chart to see the cost details for each facility type.
Additional Cost	You can see the additional expenses such as backup and restore, high availability, management, licensing, VMware software licensing.

You can select a data center to view the information specific to the data center.

Editing Cost Drivers

You can manually edit monthly cost of all the eight expense types from the current month onwards.

The configuration used for cost drivers determines how vRealize Operations Manager calculates and displays the cost.

Editing Server Hardware

You can view, add, edit, or delete the cost of each server group, based on their configuration and the purchase date of a batch server running in your cloud environment. After you update the server hardware cost, cost drivers updates the total monthly cost and average monthly cost for each server group. You can also see the details about the list of storage arrays that the EMC SRM identifies in the server.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Server Hardware**.
- 3 Click any server from the list of **Server Group Description**.

The cost drivers groups all server hardware from all data centers in your inventory based on their hardware configuration.

Category	Description
Server Group Description	Displays the name of the server in your inventory.
Number of Servers	Displays the total number of servers of any particular hardware configuration in your inventory.
Monthly Cost	Displays the average monthly cost for server. This value is calculated as a weighted average of prices of purchased and leased batches.

- 4 After selecting a server group, you can manually enter the required fields.
 - a Click **Split Batch** to enter your purchase date, cost, and type. Click split batch to enter multiple batch details of the purchase or lease of server hardware.
 - b Enter the Number of Servers and select the Purchase Type.
 - c Click **Save**.

Edit Monthly Cost of Storage

The storage hardware is categorized according to the datastore tag category. You can edit the monthly cost per storage GB for the datastores based on their storage category (using tags) and storage type (NAS, SAN, Fiber Channel or Block).

Prerequisites

To edit the cost based on storage category, you must create tags and apply them to the datastores on the vCenter Server user interface. For more information, see the VMware vSphere Documentation.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Storage**.
- 3 (Optional) Select a tag category.

Assume that you have two tag categories (for example, Profile and Tiers) with three tags in each category, you can select either Profile or Tiers from **Tag Category** to categorize the datastores based on tags.

Category	Description
Tag Category	<ul style="list-style-type: none"> ■ Category displays the tag categories for datastores and also the tags associated with the category. <p>Note If you have performed a fresh installation of vCenter Server 6.0, and not assigned tags to the datastores, cost drivers displays tag category for datastores as uncategorized.</p>
Datastores	Displays the total number of datastores for a specific category or type. You can click the datastore value to see list of datastores and its details such as monthly cost, total GB for each datastore.
Total Storage (GB)	Displays the total storage for a specific category or type.
Monthly Cost Per GB	Displays the monthly cost per GB for a specific category or type. You can edit this value for defining the monthly cost per GB for datastores.
Monthly Cost	Displays the total monthly cost for a specific category or type.

- 4 Click **Save**.

Edit Monthly Cost of License

You can edit the total operating system licensing cost and VMware license cost of your cloud environment. You can edit the license cost by either selecting the ELA charging policy or selecting the per socket value.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Drivers**.
- 2 In the Cost Drivers tab, click **License**.

The Cost drivers display all the licenses in your cloud environment.

Category	Description
Name	Displays the category of the operating system. If the operating system is not Windows or Linux, cost drivers categorize the operating system under Other Operating Systems .
VMs	Displays the number of virtual machines that are running on the specific operating system.

Category	Description
Sockets	Displays the number of sockets on which the specific operating system is running.
Charged by	Displays whether a cost is charged by socket or ELA.
Total Cost	Displays the total cost of the specific operating system.

3 Click **Save**.

Results

According to your inputs, vRealize Operations Manager calculates and displays the total cost and updates the Charged by column with the option that you have selected.

Edit Monthly Cost of Maintenance

You can edit the monthly cost of maintaining your cloud environment. Maintenance cost is categorized into hardware maintenance cost and operating system maintenance cost. Hardware maintenance cost is calculated as a percentage of the purchase cost of servers. Operating system maintenance cost is calculated as a percentage of the Windows licensing costs.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Maintenance**.
- 3 Edit the monthly maintenance cost.
 - Edit the percentage value of the hardware maintenance cost.
 - Edit the percentage value of the operating system maintenance cost.
- 4 Click **Save**.

Edit Monthly Cost of Labor

You can edit the monthly cost of labor for your cloud environment. The labor cost is combination of the total cost of the server administrator, virtual infrastructure administrator, and the operating system administrator.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Labor**.

The monthly labor cost is displayed.

Category	Description
Category	Displays the categories of labor cost, servers, virtual infrastructure, and operating system
Calculated by	Displays whether the cost is calculated hourly or monthly
Total Monthly Cost	Displays the total monthly cost of the particular category
Reference Cost	Displays the reference cost for the category from the cost drivers database

3 Click **Save**.

Results

The total monthly cost is updated. The hourly rate option or the monthly cost option that you select is updated in the **Calculated by** column.

Edit Monthly Cost of the Network

You can edit the monthly cost for each Network Interface Controller (NIC) type or can edit the total cost of all the networking expenses associated with the cloud.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Network**.
- 3 Edit the monthly cost of network.
 - Modify the values for 1 Gigabit NIC and the 10 Gigabit NIC.
 - Modify the total monthly cost of all network expenses associated with the cloud.
- 4 Click **Save**.

Results

The total monthly network expenses are updated.

Edit Monthly Cost of Facilities

For your cloud environment, you can specify the total monthly cost of facilities or edit the facilities cost for real estate, power, and cooling requirements.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Facilities**.
- 3 Edit the monthly facilities cost.
 - Modify the cost of rent or real estate per rack unit and modify the monthly cost of power and cooling per kilowatt-hour.
 - Modify the total monthly cost of facilities.
- 4 Click **Save** to save to update the changes.

Results

The monthly facilities cost is updated.

Editing Additional Costs

The additional cost lets you add any additional or extra expense that is not covered by other expenses categorized by vRealize Operations Manager. No reference value is present for this expense.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Additonal Costs**.
- 3 Enter or select the cost type for the expenses.

Note As a first time user, you must enter the cost type values manually. The values get saved and appear for all future selections.

- 4 Select the **Entity Type** and **Entity Selection**.
The **Entity Count** gets updated automatically.
- 5 Enter the **Monthly Cost per entity** .
The **Total Cost per month** gets computed automatically.
- 6 Click **Save**.

Add Application Cost

vRealize Operations Manager allows you to add the software component license cost, software update cost, or support cost of an application in your cloud environment. First, you have to define the applications in the Applications page present under the Environment tab and then you must enter the cost associated with the application in the Cost Drivers page.

Prerequisites

- Create applications in vRealize Operations Manager.

Procedure

- 1 In the menu, click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Applications**.
- 3 In the Applications page, click **+Add New Application Cost**.
- 4 Select the application name from the **Application Name** drop-down menu.
The number of VMs associated with the application is auto-populated in the **Number of VMs** column.
- 5 Enter the **Fixed Cost** for the application.
- 6 Click **Save**.

Edit Application Cost

vRealize Operations Manager allows you to edit the application cost of an application present in your cloud environment. You can only modify the cost associated with the application, as all the other attributes are predefined.

Procedure

- 1 In the menu, click Administration and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Applications**.
- 3 Click the edit icon next to the application cost you want to edit.
- 4 Modify the cost of the application.
- 5 Click **Save**.

Editing Cluster Cost Calculation Methods

You can now edit the cluster cost calculation method based on your business requirement. Earlier the cost computation was based on the actual utilization of resources. Now you can calculate the cluster utilization cost using any one of the following methods:

- Actual Utilization
- Expected Utilization across all clusters
- Expected Utilization per cluster

Procedure

- 1 In the menu, Click **Administration** and then in the left pane click **Configuration > Cost Settings**.
- 2 In the Cluster Cost tab, click **CHANGE**.
The Cluster Cost Calculation Methods dialog box is displayed.
- 3 Select any one of the Cluster Cost Calculation methods.

Option	Description
Actual Utilization	By default, the cluster cost calculation is based on the actual utilization of CPU and memory.
Expected Utilization across all clusters	You can set the fixed utilization percentages for expected CPU and memory utilization. If you select this option, the value you enter is applied across all the server clusters.
Expected Utilization per cluster	You can set the expected CPU and memory utilization percentages for each cluster by entering the value in the Expected CPU Utilization % and Expected Memory Utilization % text boxes.

- 4 Click **SAVE**.

Cluster Cost Overview

vRealize Operations Manager calculates the base rates of CPU and memory so that they can be used for virtual machine cost computation. Base rates are determined for each cluster, which are homogeneous provisioning groups. As a result, base rates might change across clusters, but are the same within a cluster.

- 1 vRealize Operations Manager first arrives at the fully loaded cost of the cluster from the cost drivers. After the cost of a cluster is determined, this cost is split into CPU and memory costs based on the industry standard cost ratios for the different models of the server.
- 2 The CPU base rate is first computed by dividing the CPU cost of the cluster by the CPU capacity of the cluster. CPU base rate is then prorated by dividing the CPU base rate by expected CPU use percentage to arrive at true base rate for charging the virtual machines.
- 3 The memory base rate is first computed by dividing the memory cost of the cluster by the memory capacity of the cluster. Memory base rate is then prorated by dividing the memory base rate by expected memory use percentage to arrive at true base rate for charging the virtual machines.
- 4 You can either provide the expected CPU and memory use or you can use the actual CPU and memory usage values.

Cluster Cost Elements	Calculation
Total Compute Cost	Total Compute Cost = (Total Infrastructure cost, which is a sum of all cost drivers) – (Storage cost) – (Direct VM cost, which is sum of OS labor, VM labor and any Windows Desktop licenses).
Expected CPU and Memory use	Expected CPU and Memory use = These percentages are arrived based on historical actual use of clusters.
Per GHz CPU base rate	Per GHz CPU base rate = (Cost attributed to CPU out of Total compute cost) / (Expected CPU Utilization * Cluster CPU Capacity in GHz).
Per GB RAM base rate	Per GB RAM base rate = (Cost attributed to RAM out of Total compute cost) / (Expected Memory Utilization * Cluster RAM Capacity in GB).
Average CPU Utilization	Average CPU Utilization = (Cost attributed to CPU utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster).
Average Memory Utilization	Average Memory Utilization = (Cost attributed to Memory utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster).
Expected CPU Utilization	The utilization percentage level of CPU that the cluster is expected to operate.
Expected Memory Utilization	The utilization percentage level of Memory that the cluster is expected to operate.

Cost Calculation Status Overview

You can check the ongoing status of manually triggered cost calculation process.

Cost calculation by default, occurs daily and whenever there is a change in the inventory or cost drivers values. You can trigger the cost calculation manually so that changes in the inventory and cost driver values reflect accordingly on the VM cost without having to wait there for any failures in the cost calculation process. It also shows default schedules time for next cost calculation process.

Migration of Cost Driver Configuration from vRealize Business for Cloud to vRealize Operations Manager

vRealize Business for Cloud supports migration of cost driver configuration from vRealize Business for Cloud to vRealize Operations Manager. You can migrate cost driver configuration from vRealize Business for Cloud 7.x or later to vRealize Operations Manager 6.7 or vRealize Operations Manager 7.0.

For more information about the migration process, see the KB article <https://kb.vmware.com/s/article/55785>.

vRealize Automation Solution

The vRealize Automation solution extends operational management capabilities of the vRealize Operations Manager platform to provide tenant-aware operational visibility of the cloud infrastructure.

The vRealize Automation solution enables you as a cloud provider to monitor the health and capacity risk of your cloud infrastructure in the context of the tenant's business groups.

You can use the vRealize Automation solution to perform some of the following key tasks:

- To gain visibility into the performance and health of the tenant's business groups that the underlying cloud infrastructure supports.
- To minimize the time taken to troubleshoot, if there is a tenant workload or an underlying infrastructure problem. The vRealize Automation solution provides visibility into the impact to performance, health, and capacity risk of the business groups because of an operational problem in the underlying cloud infrastructure layer.
- To manage the placements of VMs that are part of the clusters managed by vRealize Automation.

Supported vRealize Automation Versions

The vRealize Automation solution is supported with vRealize Automation 7.0 versions. Workload placement for day 1 operations is supported from vRealize Automation 7.3 onwards with vRealize Operations Manager 6.6 and above. Workload placement for day 2 operations is supported from vRealize Automation 7.5 onwards with vRealize Operations Manager 7.0 and above.

If you upgrade from a previous version to vRealize Operations Manager 7.0, that has the vRealize Automation Management Pack 4.0 installed, the following behavior is observed:

- vRealize Automation Management Pack 4.0 is upgraded to 7.0.

Object Types and Relationships

The vRealize Automation solution brings in cloud constructs and their relationships from vRealize Automation into vRealize Operations Manager for operational analysis.

You can use the following items in the virtual infrastructure as object types in vRealize Operations Manager.

- Tenant
- Reservation
- Business Group
- Deployment
- Blueprint
- Managed Resources
- Reservation Policy
- Virtual Machine
- Datastore
- vRealize Automation World
- vRealize Automation Management Pack Instance

Objects types in an enterprise environment are related to other objects types in that environment. Object types are either part of a larger object type, or they contain smaller component objects, or both. When you select a parent object type, vRealize Operations Manager shows any related child objects types.

Table 1-24. Relationship Model

Relationship View	Parent-Child Relationship Between Objects
Application View	Tenant > Deployment > Virtual Machine
Infrastructure View	Tenant > Business Group > Reservation > Cluster and Datastore
Blueprint View	Tenant > Business Group > Blueprint > Deployment > Virtual Machine
Deployment View	Tenant > Deployment > Virtual Machine
Reservation Policy View	Reservation Policy > Reservation > Cluster

vRealize Automation Workload Placement

You can enable workload placement when you add vRealize Operations Manager 6.6 as an endpoint in vRealize Automation 7.3. You cannot enable workload placement by adding a version of vRealize Operations Manager that is previous to version 6.6, as an endpoint in vRealize Automation 7.3.

To add vRealize Operations Manager as an endpoint in vRealize Automation 7.3, complete the following steps.

Procedure

- 1 Log in to vRealize Automation as a tenant user.
- 2 Select **Infrastructure > Endpoint > Endpoints**.
- 3 Select **New > Management > vRealize Operations Manager**.
- 4 Enter the general information for the vRealize Operations Manager endpoint.
- 5 Click **OK**.

Port Information

In environments where strict firewalls are in place, specific ports must be open for the vRealize Automation solution to retrieve data from vRealize Operations Manager.

- vRealize Automation CAFÉ Appliance/VIP URL on port 443
- vRealize Automation IAAS URL on port 443
- vRealize Automation SSO URL on port 7444

Note The vRealize Automation solution supports only vCenter objects used and managed by vRealize Automation. No other object kinds such as AWS or Openstack resources are supported at this time.

Security Guidelines

Solutions in vRealize Operations Manager execute independently. They execute within a common runtime environment within the vRealize Operations Manager collector host.

Java language security protects the adapters from interference with other adapters. All adapters execute within the common JRE process trust zone. You must only load and use adapters that you obtain from a publisher you trust and only after you verify the adapter's code integrity before loading into vRealize Operations Manager.

Even though adapters execute independently, they can make configuration changes to the collector host or Java runtime environment that may affect the security of other adapters. For example, at installation time an adapter can modify the list of trusted certificates. During execution an adapter can change the TLS/SSL certificate validation scheme and thereby change how other adapters validate certificates. The vRealize Operations Manager system and collector hosts do not isolate adapters beyond the natural isolation provided by Java execution. The system trusts all adapters equally.

Adapters are responsible for their own data security. When they collect data or make configuration changes to data sources, each adapter provides its own mechanisms and guarantees with regard to the confidentiality, integrity, and authenticity of the collected data.

The vRealize Automation solution enforces certificate checks when communicating with the vRealize Automation servers. These certificates are presented when the user clicks the **Test** button on the Adapter Instance setup page. Once these certificates are accepted by the user, they will be associated with that adapter instance. Any communication to the vRealize Automation servers will ensure that the certificates presented by the servers match the ones accepted by the user.

Configuring vRealize Automation

You can configure an instance of the vRealize Automation from which you are collecting data.

Prerequisites

- The super user must have the following privileges:
 - Infrastructure administrator rights for all tenants.
 - Infrastructure architect rights for all tenants.
 - Tenant administrator rights for all tenants.
 - Software architect roles for all tenants.
 - Fabric group administrator rights for all fabric groups, in all tenants.
- Configure the vCenter adapter instance for the same vCenter that is added as an endpoint in the vRealize Automation system.
- Use only DNS names and not IP addresses when you configure the vRealize Automation solution in a vRealize Automation distributed setup. Add host file entries on all vRealize Operations Manager nodes in the /etc/hosts location if the DNS is not reachable using vRealize Operations Manager.
- The super user account must be created for all the tenants by using an identical user name and password with the required permissions for successful data collection.

Procedure

- 1 In the menu, click **Administration**, and then from the left pane click **Solutions**.
- 2 Select VMware vRealize Automation and click the **Configure** icon.
- 3 Configure the solution.

Option	Description
Display Name	The name for the adapter instance.
Description	(Optional) The description of the adapter instance.
vRealize Automation Appliance URL	<p>The URL of the vRealize Automation CAFÉ appliance from which you are collecting data. Enter the host name, https://HostName, or the IP address, https://IP.</p> <p>If there is a load balancer for the CAFÉ appliances, the URL must have HostName or IP address of the load balancer in the format https://HostName or https://IP.</p>

Option	Description
Credential	<p>To add the credentials to access the vRealize Automation environment, click the plus sign.</p> <ul style="list-style-type: none"> ■ Credential name. The name by which you are identifying the configured credentials. ■ SysAdmin Username. The user name of the vRealize Automation system administrator. ■ SysAdmin Password. The password of the vRealize Automation system administrator. ■ SuperUser Username. The user name of the vRealize Automation super user. Create a user in vRealize Automation with specific privileges mentioned in the following note. ■ SuperUser Password. The password of the vRealize Automation super user.
Advanced Settings	To configure the advanced settings, click the drop-down menu.
Collectors/Groups	<p>The collector on which the vRealize Automation solution runs.</p> <ul style="list-style-type: none"> ■ For one collector instance, select Automatically select collector. ■ For multiple collectors, to distribute the workload and optimize performance, select the collector to manage the adapter process for this instance.
Tenants	<p>Collects data for specific tenants associated with vRealize Automation. To collect data, configure the tenants in the following manner:</p> <ul style="list-style-type: none"> ■ * (by default). Data is collected for all tenants. <p>Note</p> <ul style="list-style-type: none"> ■ Tenant test is attempted for the first two tenants that are sorted based on alphabetical order. If some tenants do not have the required privileges, then the vRealize Automation solution continues to collect data for the other tenants. Failure in collecting data for a tenant that does not have the required privileges is logged in the adapter . Log file. ■ If any of the tenants do not have the required privileges, data is not collected for that tenant. <ul style="list-style-type: none"> ■ Comma separated list. Data is collected for the specific tenants that are listed and separated by comma. ■ !. Data is collected for all tenants except the ones listed after !.
vRealize Automation Endpoint Monitoring	<ul style="list-style-type: none"> ■ Enabled: Collects and monitors data for all the vRealize Automation object types with the compute clusters under managed resources. ■ Disabled: Collects and monitors data for only the reservation object type with the compute clusters under managed resources.
vRealize Automation Enabled Intelligent Placement	Default is On . Allows vRealize Automation to manage the placements of VMs that are part of the clusters managed by vRealize Automation. This mode is always On and used for work-load placement (WLP).
Enable vRealize Automation system health monitoring	Enable or disable health monitoring of the vRealize Automation system components. For example, Cafe and IAAS.

Option	Description
vRealize Automation VA FQDN	<p>The vRealize Automation VA IP or FQDN details are required when the vRealize Automation system is HA enabled and runs behind a load balancer for component discovery.</p> <p>Enter these details only when you enable vRealize Automation system health monitoring.</p>
vRealize Automation adapter collection interval (minutes)	<p>The time interval between data collections by the vRealize Automation solution.</p> <p>Default is 15 minutes. You can increase or decrease the amount of time between data collections. It is recommended that you do not change this value in large-scale environments.</p> <p>To change this value to less than 5 minutes, you must change the collection interval value in the adapter.</p>
Tenant resource collection interval (minutes)	<p>The time interval between the data collected by the tenants in the vRealize Automation solution.</p> <p>Default is 240 minutes. You can increase or decrease the amount of time between data collections. It is recommended that you do not change this value in large-scale environments.</p> <p>To change this value to less than 5 minutes, you must change the collection interval value in the adapter.</p>
Business group resource collection interval (minutes)	<p>The time interval between the data collected by the business groups in the vRealize Automation.</p> <p>Default is 60 minutes. You can increase or decrease the amount of time between data collections. It is recommended that you do not change this value in large-scale environments.</p> <p>To change this value to less than 5 minutes, you must change the collection interval value in the adapter.</p>
Blueprint resource collection interval (minutes)	<p>The time interval between the data collected by the blueprints in the vRealize Automation solution.</p> <p>Default is 60 minutes. You can increase or decrease the amount of time between data collections. It is recommended that you do not change this value in large-scale environments.</p> <p>To change this value to less than 5 minutes, you must change the collection interval value in the adapter.</p>
Autodiscovery	<p>Discover objects automatically.</p> <ul style="list-style-type: none"> ■ To set automatic discovery for objects, select True. ■ To set the automatic discovery off, select False.

- 4 Click **Test Connection** to validate the connection.

If one of the tenant connections is successful, Test Connection is successful.

- 5 Click **Save Settings**.

Configuration Properties

In large-scale environments, multiple simultaneous API calls might cause performance problems in vRealize Automation. When an adapter sends multiple parallel requests to WAPI in particular, it

severely impacts the database. Configuration properties are used to configure the settings with appropriate values.

Table 1-25. Configuration Properties

Property Name	Description	Default Value
wapiCollectionMaxSeconds	The upper limit for the amount of time that the adapter needs to try and retrieve the data from API calls. This property must be increased in large-scale environments, in addition to increasing the adapter's collection time interval.	60 (1 minute)
wapiThreadCount	The number of threads that are querying WAPI at a time. This property might be increased or decreased based on speed or performance requirements.	2
querySuiteAPIPageSize	The number of the items to fetch in a suite API call.	100
queryVraAPIPageSize	The number of the items to fetch in a single CAFE query.	100
		<p>Note It is recommended that you keep the maximum value as 100.</p> <p>Refer to the sizing guidelines for large-scale environment guidelines: Sizing Guidelines</p>

Alert Definitions

Alert definitions are combinations of symptoms and recommendations that identify problem areas in your environment and generate alerts on which you can act. Symptom and alert definitions are defined for vRealize Automation objects. The alerts are population-based alerts based on the risk or health of a certain percentage of child objects.

The health and risk thresholds are as follows:

Health

- When 25%-50% of the child objects have health issues, the parent object triggers an alert with a Warning health level.
- When 50%-75% of the child objects have health issues, the parent object triggers an alert with an Immediate health level.
- When 75%-100% of the child objects have health issues, the parent object triggers an alert with a Critical health level.

Risk

- When 25%-50% of the child objects have risk issues, the parent object triggers an alert with a Warning risk level.

- When 50%-75% of the child objects have risk issues, the parent object triggers an alert with an Immediate risk level.
- When 75%-100% of the child objects have risk issues, the parent object triggers an alert with a Critical risk level.

vSAN

You can make vSAN operational in a production environment by using dashboards to evaluate, manage, and optimize the performance of vSAN objects and vSAN-enabled objects in your vCenter Server system.

vSAN extends the following features:

- Discovers vSAN disk groups in a vSAN datastore.
- Identifies the vSAN-enabled cluster compute resource, host system, and datastore objects in a vCenter Server system.
- Automatically adds related vCenter Server components that are in the monitoring state.

Configure a vSAN Adapter Instance

When configuring an adapter instance for vSAN, you add credentials for a vCenter Server.

Prerequisites

Only vCenter Server systems that are configured for both the vCenter adapter and the vSAN adapter appear in the inventory tree under the vSAN and Storage Devices. Verify that the vCenter Server that you use to configure the vSAN adapter instance is also configured as a vCenter adapter instance for the VMware vSphere® solution. If not, add a vCenter adapter instance for that vCenter Server.

You must open port 5989 between the host and any vRealize Operations Manager node on which the vSAN adapter resides. This is applicable when the vSAN version in vSphere is 6.6 or lower.

Procedure

- 1 In the vCenter Server text box, enter the FQDN or IP address of the vCenter Server instance to which you are connecting.

The vCenter Server FQDN or IP address must be reachable from all nodes in the vRealize Operations Manager cluster.

- 2 To add credentials on the Manage Solution page, click the plus sign.
 - a In the Credential name text box, enter the name by which you are identifying the configured credentials.
 - b Type the User name and Password for the vCenter Server instance.
 - c Click **OK**.

You configured credentials to connect to a vCenter Server instance.

3 Click **Advanced Settings.**

4 For Collection Interval, select a value of five minutes or greater.

The vSAN adapter collects Health Check Service and Performance Service metrics from vSphere objects. The Health Check Service interval is configured in the vSphere interface and is 60 minutes by default. If the Health Check Service interval is 60 minutes and the vSAN adapter collection interval is 5 minutes, the vSAN adapter reports:

- Most recent Performance Service metrics for each cycle.

5 Click **Test Connection to validate the connection with your vCenter Server instance.**

6 Accept the vCenter Server security certificate.

7 Click **Save Settings.**

Results

The adapter is added to the Adapter Instance list and is active.

What to do next

To verify that the adapter is configured and collecting data from vSAN objects, wait a few collection cycles, then view application-related data.

- Inventory Explorer. Verify that all the objects related to the vSAN instance are listed. Objects should be in the collecting state and receiving data.
- Dashboards. Verify that vSAN Capacity Overview, Migrate to vSAN, vSAN Operations Overview, and Troubleshoot vSAN, are added to the default dashboards.
- Under **Environment > vSAN and Storage Devices**, verify that the vSAN hierarchy includes the following related vCenter Server system objects:
 - vSAN World
 - Cache Disk
 - Capacity Disk
 - vSAN-enabled vCenter Server clusters
 - vSAN Fault Domains (optional)
 - vSAN-enabled Hosts
 - vSAN Datastores
 - vSAN Disk Groups
 - vSAN Datastore related VMs
 - vSAN Witness Hosts (optional)

Verify that the Adapter Instance is Connected and Collecting Data

You configured an adapter instance of vSAN with credentials for a vCenter Server. Now you want to verify that your adapter instance can retrieve information from vSAN objects in your environment.

To view the object types, in the menu, click **Administration > Configuration > Inventory Explorer > Adapter Instances > vSAN Adapter Instance > <User_Created_Instance>**.

Table 1-26. Object Types that vSAN Discovers

Object Type	Description
vSAN Adapter Instance	The vRealize Operations Management Pack for vSAN instance.
vSAN Cluster	vSAN clusters in your data center.
vSAN Datastore	vSAN datastores in your data center.
vSAN Disk Group	A collection of SSDs and magnetic disks used by vSAN.
vSAN Fault Domain	A tag for a fault domain in your data center.
vSAN Host	vSAN hosts in your data center.
vSAN Witness Host	A tag for a witness host of a stretched cluster, if the stretched cluster feature is enabled on the vSAN cluster.
vSAN World	A vSAN World is a group parent resource for all vSAN adapter instances. vSAN World displays aggregated data of all adapter instances and a single root object of the entire vSAN hierarchy.
Cache Disk	A local physical device on a host used for storing VM files in vSAN.
Capacity Disk	A local physical device on a host used for read or write caching in vSAN

The vSAN adapter also monitors the following objects discovered by the VMware vSphere adapter.

- Cluster Compute Resources
- Host System
- Datastore

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Configuration > Inventory Explorer**.
- 2 In the list of tags, expand **Adapter Instances** and expand **vSAN Adapter Instance**.
- 3 Select the adapter instance name to display the list of objects discovered by your adapter instance.

- Slide the display bar to the right to view the object status.

Object Status	Description
Collection State	If green, the object is connected.
Collection Status	If green, the adapter is retrieving data from the object.

- Deselect the adapter instance name and expand the **Object Types** tag.

Each Object Type name appears with the number of objects of that type in your environment.

What to do next

If objects are missing or not transmitting data, check to confirm that the object is connected. Then check for related alerts.

To ensure that the vSAN adapter can collect all performance data, the Virtual SAN performance service must be enabled in vSphere. For instructions on how to enable the service, see [Turn on Virtual SAN Performance Service in the VMware Virtual SAN documentation](#).

If the Virtual SAN performance service is disabled or experiencing issues, an alert is triggered for the vSAN adapter instance and the following errors appear in the adapter logs.

```
ERROR com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- Failed to collect performance metrics for Disk Group
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- vSAN Performance Service might be turned OFF.
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- (vim.fault.NotFound)
{
  faultCause = null,
  faultMessage = (vmodl.LocalizableMessage)
    [
      com.vmware.vim.binding.impl.vmodl.LocalizableMessageImpl@98e1294
    ]
}
```

Installing Optional Solutions in vRealize Operations Manager

You can extend the monitoring capabilities of vRealize Operations Manager by installing optional solutions from VMware or third parties.

VMware solutions include adapters for Storage Devices, Log Insight, NSX for vSphere, Network Devices, and VCM. Third-party solutions include AWS, SCOM, EMC Smarts, and many others. To download software and documentation for optional solutions, visit the VMware Solution Exchange at <https://marketplace.vmware.com/vsx/>.

Solutions can include dashboards, reports, alerts and other content, and adapters. Adapters are how vRealize Operations Manager manages communication and integration with other products, applications, and functions. When a management pack is installed and the solution adapters are configured, you can use the vRealize Operations Manager analytics and alerting tools to manage the objects in your environment.

If you upgrade from an earlier version of vRealize Operations Manager, your management pack files are copied to the `/usr/lib/vmware-vcops/user/plugins/.backup` file in a folder with the date and time as the folder name. Before migrating your data to your new vRealize Operations Manager instance, you must configure the adapter instances again. If you have customized the adapter, your adapter customizations are not included in the migration, and you must reconfigure the customizations.

If you update a management pack in vRealize Operations Manager to a newer version, and you have customized the adapter, your adapter customizations are not included in the upgrade, and you must reconfigure them.

Managing Solution Credentials

Credentials are the user accounts that vRealize Operations Manager uses to enable one or more solutions and associated adapters, and to establish communication with the target data sources. The credentials are supplied when you configure each adapter. You can add or modify the credential settings outside the adapter configuration process to accommodate changes to your environment.

For example, if you are modifying credentials to accommodate changes based on your password policy, the adapters configured with these credentials begin using the new user name and password to communicate between vRealize Operations Manager and the target system.

Another use of credential management is to remove misconfigured credentials. If you delete valid credentials that were in active use by an adapter, you disable the communication between the two systems.

If you need to change the configured credential to accommodate changes in your environment, you can edit the credential settings without being required to configure a new adapter instance for the target system. You can edit credential settings by clicking **Administration** on the menu, and then clicking **Credentials**.

Any adapter credential you add is shared with other adapter administrators and vRealize Operations Manager collector hosts. Other administrators might use these credentials to configure a new adapter instance or to move an adapter instance to a new host.

Manage Credentials

To configure or reconfigure credentials that you use to enable an adapter instance, you must provide the collection configuration settings, for example, user name and password, that are valid on the target system. You can also modify the connection settings for an existing credential instance.

Where You Manage Credentials

On the menu, click **Administration** and in the left pane click **Management > Credentials**.

Manage Credentials Options

The Manage Credentials dialog box is used to add new or modifies existing adapter credentials. The dialog box varies depending on the type of adapter and whether you are adding or editing. The following options describe the basic options. Depending on the solution, the options other than the basic ones vary.

Caution Any adapter credentials you add are shared with other adapter administrators and vRealize Operations Manager collector hosts. Other administrators might use these credentials to configure a new adapter instance or to move an adapter instance to a new host.

Table 1-27. Manage Credential Add or Edit Options

Option	Description
Adapter Type	Adapter type for which you are configuring the credentials.
Credential Kind	Credentials associated with the adapter. The combination of adapter and credential type affects the additional configuration options.
Credential Name	Descriptive name by which you are managing the credentials.
User Name	User account credentials that are used in the adapter configuration to connect vRealize Operations Manager to the target system.
Password	Password for the provided credentials.

Managing Collector Groups

vRealize Operations Manager uses collectors to manage adapter processes such as gathering metrics from objects. You can select a collector or a collector group when configuring an adapter instance.

If there are remote collectors in your environment, you can create a collector group, and add remote collectors to the group. When you assign an adapter to a collector group, the adapter can use any collector in the group. Use collector groups to achieve adapter resiliency in cases where the collector experiences network interruption or becomes unavailable. If this occurs, and the collector is part of a group, the total workload is redistributed among all the collectors in the group, reducing the workload on each collector.

Configuring Alerts and Actions

2

In VMware vRealize Operations Manager, alerts and actions play key roles in monitoring the objects.

This chapter includes the following topics:

- [Types of Alerts](#)
- [Configuring Alerts](#)
- [Configuring Actions](#)

Types of Alerts

Different types of alerts are triggered on a certain object.

The alerts are of three types:

- Health Alerts
- Risk Alerts
- Efficiency Alerts

Configuring Alerts

Whenever there is a problem in the environment, the alerts are generated. You can create the alert definitions so that the generated alerts tell you about the problems in the monitored environment.

Defining Alerts in vRealize Operations Manager

An alert definition comprises one or more symptom definitions, and the alert definition is associated with a set of recommendations and actions that help you resolve the problem. Alert definitions include triggering symptom definitions and actionable recommendations. You create the alert definitions so that the generated alerts tell you about problems in the monitored environment. You can then respond to the alerts with effective solutions that are provided in the recommendations.

Predefined alerts are provided in vRealize Operations Manager as part of your configured adapters. You can add or modify alert definitions to reflect the needs of your environment.

Symptoms in Alert Definitions

Symptom definitions evaluate conditions in your environment that, if the conditions become true, trigger a symptom and can result in a generated alert. You can add symptom definitions that are based on metrics or super metrics, properties, message events, fault events, or metric events. You can create a symptom definition as you create an alert definition or as an individual item in the appropriate symptom definition list.

When you add a symptom definition to an alert definition, it becomes a part of a symptom set. A symptom set is the combination of the defined symptom with the argument that determines when the symptom condition becomes true.

A symptom set combines one or more symptom definitions by applying an Any or All condition, and allows you to choose the presence or absence of a particular symptom. If the symptom set pertains to related objects rather than to Self, you can apply a population clause to identify a percentage or a specific count of related objects that exhibit the included symptom definitions.

An alert definition comprises one or more symptom sets. If an alert definition requires all of the symptom sets to be triggered before generating an alert, and only one symptom set is triggered, an alert is not generated. If the alert definition requires only one of several symptom sets to be triggered, then the alert is generated even though the other symptom sets were not triggered.

Recommendations in Alert Definitions

Recommendations are the remediation options that you provide to your users to resolve the problems that the generated alert indicates.

When you add an alert definition that indicates a problem with objects in your monitored environment, add a relevant recommendation. Recommendations can be instructions to your users, links to other information or instruction sources, or vRealize Operations Manager actions that run on the target systems.

Modifying Alert Definitions

If you modify the alert impact type of an alert definition, any alerts that are already generated will have the previous impact level. Any new alerts will be at the new impact level. If you want to reset all the generated alerts to the new level, cancel the old alerts. If they are generated after cancellation, they will have the new impact level.

Defining Symptoms for Alerts

Symptoms are conditions that indicate problems in your environment. You define symptoms that you add to alert definitions so that you know when a problem occurs with your monitored objects.

As data is collected from your monitored objects, the data is compared to the defined symptom condition. If the condition is true, then the symptom is triggered.

You can define symptoms based on metrics and super metrics, properties, message events, fault events, and metric events.

Defined symptoms in your environment are managed in the Symptom Definitions. When the symptoms that are added to an alert definition are triggered, they contribute to a generated alert.

Define Symptoms to Cover All Possible Severities and Conditions

Use a series of symptoms to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of `Warning` while `Volume reached capacity limit` might have a severity level of `Critical`. The first symptom is not an immediate threat. The second symptom is an immediate threat.

About Metrics and Super Metrics Symptoms

Metric and super metric symptoms are based on the operational or performance values that vRealize Operations Manager collects from target objects in your environment. You can configure the symptoms to evaluate static thresholds or dynamic thresholds.

You define symptoms based on metrics so that you can create alert definitions that let you know when the performance of an object in your environment is adversely affected.

Static Thresholds

Metric symptoms that are based on a static threshold compare the currently collected metric value against the fixed value you configure in the symptom definition.

For example, you can configure a static metric symptom where, when the virtual machine CPU workload is greater than 90, a critical symptom is triggered.

Dynamic Thresholds

Metric symptoms that are based on dynamic thresholds compare the currently collected metric value against the trend identified by vRealize Operations Manager, evaluating whether the current value is above, below, or generally outside the trend.

For example, you can configure a dynamic metric symptom where, when the virtual machine CPU workload is above the trended normal value, a critical symptom is triggered.

Property Symptoms

Property symptoms are based on the configuration properties that vRealize Operations Manager collects from the target objects in your environment.

You define symptoms based on properties so that you can create alert definitions that let you know when changes to properties on your monitored objects can affect the behavior of the objects in your environment.

Message Event Symptoms

Message event symptoms are based on events received as messages from a component of vRealize Operations Manager or from an external monitored system through the system's REST API. You define symptoms based on message events to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

The adapters for the external monitored systems and the REST API are inbound channels for collecting events from external sources. Adapters and the REST server both run in the vRealize Operations Manager system. The external system sends the messages, and vRealize Operations Manager collects them.

You can create message event symptoms for the supported event types. The following list is of supported event types with example events.

- **System Performance Degradation.** This message event type corresponds to the `EVENT_CLASS_SYSTEM` and `EVENT_SUBCLASS_PERFORM_DEGRADATION` type and subtype in the vRealize Operations Manager API SDK.
- **Change.** The VMware adapter sends a change event when the CPU limit for a virtual machine is changed from unlimited to 2 GHz. You can create a symptom to detect CPU contention issues as a result of this configuration change. This message event type corresponds to the `EVENT_CLASS_CHANGE` and `EVENT_SUBCLASS_CHANGE` type and subtype in the vRealize Operations Manager API SDK.
- **Environment Down.** The vRealize Operations Manager adapter sends an environment down event when the collector component is not communicating with the other components. You can create a symptom that is used for internal health monitoring. This message event type corresponds to the `EVENT_CLASS_ENVIRONMENT` and `EVENT_SUBCLASS_DOWN` type and subtype in the vRealize Operations Manager API SDK.
- **Notification.** This message event type corresponds to the `EVENT_CLASS_NOTIFICATION` and `EVENT_SUBCLASS_EXTEVENT` type and subtype in the vRealize Operations Manager API SDK.

Fault Symptoms

Fault symptoms are based on events published by monitored systems. vRealize Operations Manager correlates a subset of these events and delivers them as faults. Faults are intended to signify events in the monitored systems that affect the availability of objects in your environment. You define symptoms based on faults to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

You can create fault symptoms for the supported published faults. Some object types have multiple fault definitions from which to choose, while others have no fault definitions.

If the adapter published fault definitions for an object type, you can select one or more fault events for a given fault while you define the symptom. The symptom is triggered if the fault is active because of any of the chosen events. If you do not select a fault event, the symptom is triggered if the fault is active because of a fault event.

Metric Event Symptoms

Metric event symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not vRealize Operations Manager.

Metric event symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that vRealize Operations Manager is actively monitoring.

The metric event thresholds, which determine whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system, represent the type and subtype combination that is specified in the incoming metric event.

- Above Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_ABOVE` defined in the vRealize Operations Manager API SDK.
- Below Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_BELOW` defined in the vRealize Operations Manager API SDK.
- Equal Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_EQUAL` defined in the vRealize Operations Manager API SDK.
- Not Equal Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_NOT_EQUAL` defined in the vRealize Operations Manager API SDK.

Understanding Negative Symptoms for vRealize Operations Manager Alerts

Alert symptoms are conditions that indicate problems in your environment. When you define an alert, you include symptoms that generate the alert when they become true in your environment. Negative symptoms are based on the absence of the symptom condition. If the symptom is not true, the symptom is triggered.

To use the absence of the symptom condition in an alert definition, you negate the symptom in the symptom set.

All defined symptoms have a configured criticality. However, if you negate a symptom in an alert definition, it does not have an associated criticality when the alert is generated.

All symptom definitions have a configured criticality. If the symptom is triggered because the condition is true, the symptom criticality will be the same as the configured criticality. However, if you negate a symptom in an alert definition and the negation is true, it does not have an associated criticality.

When negative symptoms are triggered and an alert is generated, the effect on the criticality of the alert depends on how the alert definition is configured.

The following table provides examples of the effect negative symptoms have on generated alerts.

Table 2-1. Negative Symptoms Effect on Generated Alert Criticality

Alert Definition Criticality	Negative Symptom Configured Criticality	Standard Symptom Configured Criticality	Alert Criticality When Triggered
Warning	One Critical Symptom	One Immediate Symptom	Warning. The alert criticality is based on the defined alert criticality.
Symptom Based	One Critical Symptom	One Warning Symptom	Warning. The negative symptom has no associated criticality and the criticality of the standard symptom determines the criticality of the generated alert.
Symptom Based	One Critical Symptom	No standard symptom included	Info. Because an alert must have a criticality and the negative alert does not have an associated criticality, the generated alert has a criticality of Info, which is the lowest possible criticality level.

Defining Recommendations for Alert Definitions

Recommendations are instructions to your users who are responsible for responding to alerts. You add recommendations to vRealize Operations Manager alerts so that your users can maintain the objects in your environment at the required levels of performance.

Recommendations provide your network engineers or virtual infrastructure administrators with information to resolve alerts.

Depending on the knowledge level of your users, you can provide more or less information, including the following options, in any combination.

- One line of instruction.
- Steps to resolve the alert on the target object.
- Hyperlink to a Web site, runbook, wiki, or other source.
- Action that makes a change on the target object.

When you define an alert, provide as many relevant action recommendations as possible. If more than one recommendation is available, arrange them in priority order so that the solution with the lowest effect and highest effectiveness is listed first. If no action recommendation is available, add text recommendations. Be as precise as possible when describing what the administrator should do to fix the alert.

Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for vRealize Operations Manager to alert you. When the alert is triggered on your host system, vRealize Operations Manager alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have vRealize Operations Manager notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, select **Alert Settings > Alert Definitions**.
- 2 Enter **capacity** in the search text box.

Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.

- 3 Click the plus sign to create a new capacity alert definition for your host systems.

- a In the alert definition workspace, for the Name and Description, enter **Hosts – Alert on Capacity Exceeded**.
- b For the Base Object Type, select **vCenter Adapter > Host System**
- c For the Alert Impact, select the following options.

Option	Selection
Impact	Select Risk .
Criticality	Select Immediate .
Alert Type and Subtype	Select Application : Capacity .
Wait Cycle	Select 1 .
Cancel Cycle	Select 1 .

- d For Add Symptom Definitions, select the following options.

Option	Selection
Defined On	Select Self .
Symptom Definition Type	Select Metric / Supermetric .
Quick filter (Name)	Enter capacity .

- e From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the right pane.

In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.

- f For Add Recommendations, enter **virtual machine** in the quick filter text box.
- g Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the right pane.

This recommendation is set to Priority 1.

- 4 Click **Save** to save the alert definition.

Your new alert appears in the list of alert definitions.

Results

You have added an alert definition to have vRealize Operations Manager alert you when the capacity of your host systems begins to run out.

Alert Definition Best Practices

As you create alert definitions for your environment, apply consistent best practices so that you optimize alert behavior for your monitored objects.

Alert Definitions Naming and Description

The alert definition name is the short name that appears in the following places:

- In data grids when alerts are generated
- In outbound alert notifications, including the email notifications that are sent when outbound alerts and notifications are configured in your environment

Ensure that you provide an informative name that clearly states the reported problem. Your users can evaluate alerts based on the alert definition name.

The alert definition description is the text that appears in the alert definition details and the outbound alerts. Ensure that you provide a useful description that helps your users understand the problem that generated the alert.

Wait and Cancel Cycle

The wait cycle setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition goes into effect after the wait cycle for the symptom definition results in a triggered symptom. In most alert definitions you configure the sensitivity at the symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that the alert is immediately generated after all of the symptoms are triggered at the desired symptom sensitivity level.

The cancel cycle setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition goes into effect after the cancel cycle for the symptom definition results in a cancelled symptom. In most definitions you configure the sensitivity at the symptom level and configure the cancel cycle of alert definition to 1. This configuration ensures that the alert is immediately cancelled after all of the symptoms conditions disappear after the desired symptom cancel cycle.

Create Alert Definitions to Generate the Fewest Alerts

You can control the size of your alert list and make it easier to manage. When an alert is about a general problem that can be triggered on a large number of objects, configure its definition so that the alert is generated on a higher level object in the hierarchy rather than on individual objects.

As you add symptoms to your alert definition, do not overcrowd a single alert definition with secondary symptoms. Keep the combination of symptoms as simple and straightforward as possible.

You can also use a series of symptom definitions to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of `Warning` while `Volume reached capacity limit` might have a severity level of `Critical`. The first symptom is not an immediate threat, but the second one is an immediate threat. You can then include the `Warning` and `Critical` symptom definitions in a single alert definition with an `Any` condition and set the alert criticality to be `Symptom Based`. These settings cause the alert to be generated with the right criticality if either of the symptoms is triggered.

Avoid Overlapping and Gaps Between Alerts

Overlaps result in two or more alerts being generated for the same underlying condition. Gaps occur when an unresolved alert with lower severity is canceled, but a related alert with a higher severity cannot be triggered.

A gap occurs in a situation where the value is $\leq 50\%$ in one alert definition and $\geq 75\%$ in a second alert definition. The gap occurs because when the percentage of volumes with high use falls between 50 percent and 75 percent, the first problem cancels but the second does not generate an alert. This situation is problematic because no alert definitions are active to cover the gap.

Actionable Recommendations

If you provide text instructions to your users that help them resolve a problem identified by an alert definition, precisely describe how the engineer or administrator should fix the problem to resolve the alert.

To support the instructions, add a link to a wiki, runbook, or other sources of information, and add actions that you run from vRealize Operations Manager on the target systems.

Creating and Managing vRealize Operations Manager Alert Notifications

When alerts are generated in vRealize Operations Manager, they appear in the alert details and object details, but you can also configure vRealize Operations Manager to send your alerts to outside applications using one or more outbound alert options.

You configure notification options to specify which alerts are sent out for the `Standard Email`, `REST`, `SNMP`, and `Log File` outbound alert plug-ins. For the other plug-in types, all the alerts are sent when the target outbound alert plug-in is enabled.

The most common outbound alert plug-in is the Standard Email plug-in. You configure the Standard Email plug-in to send notifications to one or more users when an alert is generated that meets the criteria you specify in the notification settings.

List of Outbound Plug-Ins in vRealize Operations Manager

vRealize Operations Manager provides outbound plug-ins. This list includes the name of the plug-in and whether you can filter the outbound data based on your notification settings.

If the plug-in supports configuring notification rules, then you can filter the messages before they are sent to the target system. If the plug-in does not support notifications, all messages are sent to the target system, and you can process them in that application.

If you installed other solutions that include other plug-in options, they appear as a plug-in option with the other plug-ins.

Messages and alerts are sent only when the plug-in is enabled.

Table 2-2. Notification Support for Outbound Plug-Ins

Outbound Plug-In	Configure Notification Rules
Automated Action Plug-in	No The Automated Action plug-in is enabled by default. If automated actions stop working, check the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.
Log File Plug-In	Yes To filter the log file alerts, you can either configure the file named <code>TextFilter.xml</code> or configure the notification rules.
Smarts SAM Notification Plug-In	No
REST Notification Plug-In	Yes
Network Share Plug-In	No
Standard Email Plug-In	Yes
SNMP Trap Plug-In	Yes

Add Outbound Notification Plug-Ins in vRealize Operations Manager

You add outbound plug-in instances so that you can notify users about alerts or capture alert data outside of vRealize Operations Manager.

You can configure one or more instances of the same plug-in type if you need to direct alert information to multiple target systems.

The Automated Action plug-in is enabled by default. If automated actions stop working, check the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.

- [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email vRealize Operations Manager alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals.

- [Add a REST Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a REST Plug-In so that you can send vRealize Operations Manager alerts to another REST-enabled application where you built a REST Web service to accept these messages.

- [Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.

- [Add a Network Share Plug-In for vRealize Operations Manager Reports](#)

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports SMB version 2.0.

- [Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

- [Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts

You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email vRealize Operations Manager alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals.

Prerequisites

Ensure that you have an email user account that you can use as the connection account for the alert notifications. If you choose to require authentication, you must also know the password for this account.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click the plus sign to add a plug-in.
- 3 From the **Plug-In Type** drop-down menu, select **Standard Email Plugin**.

The dialog box expands to include your SMTP settings.

4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

5 Configure the SMTP options appropriate for your environment.

Option	Description
Use Secure Connection	Enables secure communication encryption using SSL/TLS. If you select this option, you must select a method in the Secure Connection Type drop-down menu.
Requires Authentication	Enables authentication on the email user account that you use to configure this SMTP instance. If you select this option, you must provide a password for the user account.
SMTP Host	URL or IP address of your email host server.
SMTP Port	Default port SMTP uses to connect with the server.
Secure Connection Type	Select either SSL/TLS as the communication encryption method used in your environment from the drop-down menu. You must select a connection type if you select Use Secure Connection.
User Name	Email user account that is used to connect to the email server.
Password	Password for the connection user account. A password is required if you select Requires Authentication.
Sender Email Address	Email address that appears on the notification message
Sender Name	Displayed name for the sender email address.

6 Click **Save**.

7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the standard email plug-in for outbound SMTP alerts is configured and running.

What to do next

Create notification rules that use the standard email plug-in to send a message to your users about alerts requiring their attention. See [User Scenario: Create a vRealize Operations Manager Email Alert Notification](#).

Add a REST Plug-In for vRealize Operations Manager Outbound Alerts

You add a REST Plug-In so that you can send vRealize Operations Manager alerts to another REST-enabled application where you built a REST Web service to accept these messages.

The REST Plug-In supports enabling an integration, it does not provide an integration. Depending on your target application, you might need an intermediary REST service or some other mechanism that will correlate the alert and object identifiers included in the REST alert output with the identifiers in your target application.

Determine which content type you are delivering to your target application. If you select application/json, the body of the POST or PUT calls that are sent have the following format. Sample data is included.

```
{
  "startDate":1369757346267,
  "criticality":"ALERT_CRITICALITY_LEVEL_WARNING",
  "Risk":4.0,
  "resourceId":"sample-object-uuid",
  "alertId":"sample-alert-uuid",
  "status":"ACTIVE",
  "subType":"ALERT_SUBTYPE_AVAILABILITY_PROBLEM",
  "cancelDate":1369757346267,
  "resourceKind":"sample-object-type",
  "alertName":"Invalid IP Address for connected Leaf Switch",
  "attributeKeyID":5325,
  "Efficiency":1.0,
  "adapterKind":"sample-adapter-type",
  "Health":1.0,
  "type":"ALERT_TYPE_APPLICATION_PROBLEM",
  "resourceName":"sample-object-name",
  "updateDate":1369757346267,
  "info":"sample-info"
}
```

If you select application/xml, the body of the POST or PUT calls that are sent have the following format:

```
<alert>
  <startDate>1369757346267</startDate>
  <criticality>ALERT_CRITICALITY_LEVEL_WARNING</criticality>
  <Risk>4.0</Risk>
  <resourceId>sample-object-uuid</resourceId>
  <alertId>sample-alert-uuid</alertId>
  <status>ACTIVE</status>
  <subType>ALERT_SUBTYPE_AVAILABILITY_PROBLEM</subType>
  <cancelDate>1369757346267</cancelDate>
  <resourceKind>sample-object-type</resourceKind>
  <alertName>Invalid IP Address for connected Leaf Switch</alertName>
  <attributeKeyId>5325</attributeKeyId>
  <Efficiency>1.0</Efficiency>
  <adapterKind>sample-adapter-type</adapterKind>
  <Health>1.0</Health>
  <type>ALERT_TYPE_APPLICATION_PROBLEM</type>
  <resourceName>sample-object-name</resourceName>
  <updateDate>1369757346267</updateDate>
  <info>sample-info</info>
</alert>
```

Note If the alert is triggered by a non-metric violation, the attributeKeyID is omitted from the REST output and is not sent.

If the request is processed as POST, for either JSON or XML, the Web service returns an HTTP status code of 201, which indicates the alert was successfully created at the target. If the request is processed as PUT, the HTTP status code of 202, which indicates the alert was successfully accepted at the target.

Prerequisites

Ensure that you know how and where the alerts sent using the REST plug-in are consumed and processed in your environment, and that you have the appropriate connection information available.

Procedure

- 1 In the left pane of vRealize Operations Manager, click the **Administration** icon.
- 2 Click **Outbound Settings** and click the plus sign to add a plug-in.
- 3 From the **Plug-In Type** drop-down menu, select **Rest Notification Plugin**.

The dialog box expands to include your REST settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Rest options appropriate for your environment.

Option	Description
URL	URL to which you are sending the alerts. The URL must support HTTPS. When an alert is sent to the REST Web server, the plug-in appends / {alertID} to the POST or PUT call.
User Name	User account on the target REST system.
Password	User account password.
Content Type	Specify the format for the alert output. <ul style="list-style-type: none"> ■ application/json. Alert data is transmitted using JavaScript Object Notation as human-readable text. ■ application/xml. Alert data is transmitted using XML that is human-readable and machine-readable content.
Certificate thumbprint	Thumbprint for the public certificate for your HTTPS service. Either the SHA1 or SHA256 algorithm can be used.
Connection count	Limits the number of simultaneous alerts that are sent to the target REST server. Use this number to ensure that your REST server is not overwhelmed with requests.

- 6 Click **Save**.
- 7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the REST plug-in for outbound alerts is configured and running.

What to do next

Create notification rules that use the REST plug-in to send alerts to a REST-enabled application or service in your environment. See [User Scenario: Create a vRealize Operations Manager REST Alert Notification](#).

Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts

You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.

All alerts are added to the log file. You can use other applications to filter and manage the logs.

Prerequisites

Ensure that you have write access to the file system path on the target vRealize Operations Manager nodes.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click the plus sign to add a plug-in.
- 3 From the **Plug-In Type** drop-down menu, select **Log File**.
The dialog box expands to include your log file settings.
- 4 In the **Alert Output Folder** text box, enter the folder name.
If the folder does not exist in the target location, the plug-in creates the folder in the target location. The default target location is: `/usr/lib/vmware-vcops/common/bin/`.
- 5 Click **Save**.
- 6 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the log file plug-in is configured and running.

What to do next

When the plug-in is started, the alerts are logged in the file. Verify that the log files are created in the target directory as the alerts are generated, updated, or canceled.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports SMB version 2.0.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.

- 2 From the toolbar, click the **Add** icon.

- 3 From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

The dialog box expands to include your plug-in instance settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.
Password	The password for the domain user account.
Network share root	<p>The path to the root folder where you want to save the reports. You can specify subfolders for each report when you configure the schedule publication.</p> <p>You must enter an IP address. For example, <code>\\IP_address\ShareRoot</code>. You can use the host name instead of the IP address if the host name is resolved to an IPv4 when accessed from the vRealize Operations Manager host.</p> <p>Note Verify that the root destination folder exists. If the folder is missing, the Network Share plug-in logs an error after 5 unsuccessful attempts.</p>

- 6 Click **Test** to verify the specified paths, credentials, and permissions.

The test might take up to a minute.

- 7 Click **Save**.

The outbound service for this plug-in starts automatically.

- 8 (Optional) To stop an outbound service, select an instance and click **Disable** on the toolbar.

Results

This instance of the Network Share plug-in is configured and running.

What to do next

Create a report schedule and configure it to send reports to your shared folder.

Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts

You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

You can provide filtering when you define a Notification using an SNMP Trap destination.

Prerequisites

Ensure that you have an SNMP Trap server configured in your environment, and that you know the IP address or host name, port number, and community that it uses.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click the plus sign to add a plug-in.
- 3 From the **Plug-In Type** drop-down menu, select **SNMP Trap**.
The dialog box expands to include your SNMP trap settings.
- 4 Type an **Instance Name**.
- 5 Configure the SNMP trap settings appropriate to your environment.

Option	Description
Destination Host	IP address or fully qualified domain name of the SNMP management system to which you are sending alerts.
Port	Port used to connect to the SNMP management system. Default port is 162.
Community	Text string that allows access to the statistics. SNMP Community strings are used only by devices that support SNMPv3 protocol.
Username	Username to configure SNMP trap settings in your environment. If the username is specified, SNMPv3 is considered as the protocol by the plugin. If left blank, SNMPv2c is considered as the protocol by the plugin.
Authentication Protocol	Authentication algorithms available are SHA-224, SHA-256, SHA-384, SHA-512.
Authentication Password	Authentication password.
Privacy Protocol	Privacy algorithms available are AES192, AES2564.
Privacy Password	Privacy password.

- 6 Click **Save**.

Results

This instance of the SNMP Trap plug-in is configured and running.

What to do next

When the plug-in is added, [Configuring Notifications](#) for receiving the SNMP traps.

Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts

You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

This outbound alert option is useful when you manage the same objects in Server Assurance Manager and in vRealize Operations Manager, and you added the EMC Smarts management pack and configured the solution in vRealize Operations Manager. Although you cannot filter the alerts sent to Service Assurance Manager in vRealize Operations Manager, you can configure the Smarts plug-in to send the alerts to the Smarts Open Integration server. You then configure the Open Integration server to filter the alerts from vRealize Operations Manager, and send only those that pass the filter test to the Smarts Service Assurance Manager service.

Prerequisites

- Verify that you configured the EMC Smarts solution. For documentation regarding EMC Smarts integration, see <https://solutionexchange.vmware.com/store>.
- Ensure that you have the EMC Smarts Broker and Server Assurance Manager instance host name or IP address, user name, and password.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click the plus sign to add a plug-in.
- 3 From the **Plug-In Type** drop-down menu, select **Smarts SAM Notification**.

The dialog box expands to include your Smarts settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Smarts SAM notification settings appropriate for your environment.

Option	Description
Broker	Type the host name or IP address of the EMC Smarts Broker that manages registry for the Server Assurance Manager instance to which you want the notifications sent.
Broker Username	If the Smarts broker is configured as Secure Broker, type the user name for the Broker account.
Broker Password	If the Smarts broker is configured as Secure Broker, type the password for the Broker user account.
SAM Server	Type the host name or IP address of the Server Assurance Manager server to which you are sending the notifications.

Option	Description
User Name	Type the user name for the Server Assurance Manager server instance. This account must have read and write permissions for the notifications on the Smarts server as specified in the SAM Server.
Password	Type the password for the Server Assurance Manager server account.

6 Click **Save**.

7 Modify the Smarts SAM plug-in properties file.

- a Open the properties file at: `/usr/lib/vmware-vcops/user/plugins/outbound/vcops-smartsalert-plugin/conf/plugin.properties`
- b Add the following string to the properties file: #
`sendByType=APPLICATION::AVAILABILITY,APPLICATION::PERFORMANCE,APPLICATION::CAPACITY,APPLICATION::COMPLIANCE,VIRTUALIZATION::AVAILABILITY,VIRTUALIZATION::PERFORMANCE,VIRTUALIZATION::CAPACITY,VIRTUALIZATION::COMPLIANCE,HARDWARE::AVAILABILITY,HARDWARE::PERFORMANCE,HARDWARE::CAPACITY,HARDWARE::COMPLIANCE,STORAGE::AVAILABILITY,STORAGE::PERFORMANCE,STORAGE::CAPACITY,STORAGE::COMPLIANCE,NETWORK::AVAILABILITY,NETWORK::PERFORMANCE,NETWORK::CAPACITY,NETWORK::COMPLIANCE`
- c Save the properties file.

8 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the Smarts SAM Notifications plug-in is configured and running.

What to do next

In Smarts Service Assurance Manager, configure your Notification Log Console to filter the alerts from vRealize Operations Manager. To configure the filtering for Service Assurance Manager, see the EMC Smarts Service Assurance Manager documentation.

Configuring Notifications

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager. You configure notification rules for the supported outbound alerts so that you can filter the alerts that are sent to the selected external system.

You use the notifications list to manage your rules. You then use the notification rules to limit the alerts that are sent to the external system. To use notifications, the supported outbound alert plug-ins must be added and running.

With notification rules, you can limit the data that is sent to the following external systems.

- **Standard Email.** You can create multiple notification rules for various email recipients based on one or more of the filter selections. If you add recipients but do not add filter selections, all the generated alerts are sent to the recipients.

- **REST.** You can create a rule to limit alerts that are sent to the target REST system so that you do not need to implement filtering on that target system.
- **SNMP Trap.** You can configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.
- **Log File.** You can configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes.

User Scenario: Create a vRealize Operations Manager Email Alert Notification

As a virtual infrastructure administrator, you need vRealize Operations Manager to send email notifications to your advanced network engineers when critical alerts are generated for `mmbhost` object, the host for many virtual machines that run transactional applications, where no one has yet taken ownership of the alert.

Prerequisites

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see [Create an Alert Definition for Department Objects](#).
- Ensure that at least one instance of the standard email plug-in is configured and running. See [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.
- 2 Click **Notification Settings** and click the plus sign to add a notification rule.
- 3 In the **Name** text box type a name similar to **Unclaimed Critical Alerts for mmbhost**.
- 4 In the Method area, select **Standard Email Plug-In** from the drop-down menu, and select the configured instance of the email plug-in.
- 5 Configure the email options.
 - a In the **Recipients** text box, type the email addresses of the members of your advance engineering team, separating the addresses with a semi-colon (;).
 - b To send a second notification if the alert is still active after a specified amount of time, type the number of minutes in the **Notify again** text box.
 - c Type number of notifications that are sent to users in the **Max Notifications** text box.
- 6 Configure the scope of filtering criteria.
 - a From the **Scope** drop-down menu, select **Object**.
 - b Click **Click to select Object** and type the name of the object.
In this example, type `mmbhost`.
 - c Locate and select the object in the list, and click **Select**.

- 7 Configure the Notification Trigger.
 - a From the **Notification Trigger** drop-down menu, select **Impact**.
 - b From the adjacent drop-down menu, select **Health**.
- 8 In the Criticality area, click **Critical**.
- 9 Expand the Advanced Filters and from the **Alert States** drop-down menu, select **Open**.
The Open state indicates that no engineer or administrator has taken ownership of the alert.
- 10 Click **Save**.

Results

You created a notification rule that sends an email message to the members of your advance network engineering team when any critical alerts are generated for the mmbhost object and the alert is not claimed by an engineer. This email reminds them to look at the alert, take ownership of it, and work to resolve the triggering symptoms.

What to do next

Respond to alert email notifications. See *vRealize Operations Manager User Guide*.

User Scenario: Create a vRealize Operations Manager REST Alert Notification

As a virtual infrastructure administrator, you need vRealize Operations Manager to send alerts in JSON or XML to a REST-enabled application that has REST Web service that accepts these messages. You want only alerts where the virtualization alerts that affect availability alert types go to this outside application. You can then use the provided information to initiate a remediation process in that application to address the problem indicated by the alert.

The notification configuration limits the alerts sent to the outbound alert instance to those matching the notification criteria.

Prerequisites

- Verify that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see [Create an Alert Definition for Department Objects](#).
- Verify that at least one instance of the REST plug-in is configured and running. See [Add a REST Plug-In for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.
- 2 Click **Notifications** and click the plus sign to add a notification rule.
- 3 In the **Name** text box type a name similar to **Virtualization Alerts for Availability**.
- 4 In the Method area, select **REST Plug-In** from the drop-down menu, and select the configured instance of the email plug-in.

- 5 Configure the Notification Trigger.
 - a From the **Notification Trigger** drop-down menu, select **Alert Type**.
 - b Click **Click to select Alert type/subtype** and select **Virtualization/Hypervisor Alerts Availability**.
- 6 In the Criticality area, click **Warning**.
- 7 Expand the Advanced Filters and from the **Alert Status** drop-down menu, select **New**.
The New status indicates that the alert is new to the system and not updated.
- 8 Click **Save**.

Results

You created a notification rule that sends the alert text to the target REST-enabled system. Only the alerts where the configured alert impact is Virtualization/Hypervisor Availability and where the alert is configured as a warning are sent to the target instance using the REST plug-in.

Create an Alert Definition for Department Objects

As a virtual infrastructure administrator, you are responsible for the virtual machines and hosts that the accounting department uses. You can create alerts to manage the accounting department objects.

You received several complaints from your users about delays when they are using their accounting applications. Using vRealize Operations Manager, you identified the problem as related to CPU allocations and workloads. To better manage the problem, you create an alert definition with tighter symptom parameters so that you can track the alerts and identify problems before your users encounter further problems.

Using this scenario, you create a monitoring system that monitors your accounting objects and provides timely notifications when problems occur.

Add Description and Base Object to Alert Definition

To create an alert to monitor the CPUs for the accounting department virtual machines and monitor host memory for the hosts on which they operate, you begin by describing the alert.

When you name the alert definition and define alert impact information, you specify how the information about the alert appears in vRealize Operations Manager. The base object is the object around which the alert definition is created. The symptoms can be for the base object and for related objects.

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Alert Settings > Alert Definitions**.
- 2 Click the plus sign to add a definition.

3 Type a name and description.

In this scenario, type **Acct VM CPU early warning** as the alert name, which is a quick overview of the problem. The description, which is a detailed overview, should provide information that is as useful as possible. When the alert is generated, this name and description appears in the alert list and in the notification.

4 Click **Base Object Type**.

5 From the drop-down menu, expand **vCenter Adapter** and select **Host System**.

This alert is based on host systems because you want an alert that acts as an early warning to possible CPU stress on the virtual machines used in the accounting department. By using host systems as the based object type, you can respond to the alert symptom for the virtual machines with bulk actions rather than responding to an alert for each virtual machine.

6 Click **Alert Impact** and configure the metadata for this alert definition.

a From the **Impact** drop-down menu, select **Risk**.

This alert indicates a potential problem and requires attention in the near future.

b From the **Criticality** drop-down menu, select **Immediate**.

As a Risk alert, which is indicative of a future problem, you still want to give it a high criticality so that it is ranked for correct processing. Because it is designed as an early warning, this configuration provides a built-in buffer that makes it an immediate risk rather than a critical risk.

c From the **Alert Type and Subtype** drop-down menu, expand **Virtualization/Hypervisor** and select **Performance**.

d To ensure that the alert is generated during the first collection cycle after the symptoms become true, set the **Wait Cycle** to **1**.

e To ensure that the an alert is removed as soon as the symptoms are no longer triggered, set the **Cancel Cycle** to **1**.

The alert is canceled in the next collection cycle if the symptoms are no long true.

These alert impact options help you identify and prioritize alerts as they are generated.

Results

You started an alert definition where you provided the name and description, selected host system as the base object type, and defined the data that appears when the alert generated.

What to do next

Continue in the workspace, adding symptoms to your alert definition. See [Add a Virtual Machine CPU Usage Symptom to the Alert Definition](#).

Add a Virtual Machine CPU Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add symptoms to your vRealize Operations Manager alert definition after you provide the basic descriptive information for the alert. The first symptom you add is related to CPU usage on virtual machines. You later use a policy and group to apply alert to the accounting virtual machines.

This scenario has two symptoms, one for the accounting virtual machines and one to monitor the hosts on which the virtual machines operate.

Prerequisites

Begin configuring the alert definition. See [Add Description and Base Object to Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Add Symptom Definitions** and configure the symptoms.
- 2 Begin configuring the symptom set related to virtual machines CPU usage.
 - a From the **Defined On** drop-down menu, select **Child**.
 - b From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.
 - c From the **Symptom Definition Type** drop-down menu, select **Metric / Supermetric**.
 - d Click the **Add** button to open the **Add Symptom Definition** workspace window.
- 3 Configure the virtual machine CPU usage symptom in the **Add Symptom Definition** workspace window.
 - a From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Virtual Machine**.
 The collected metrics for virtual machines appears in the list.
 - b In the metrics list **Search** text box, which searches the metric names, type **usage**.
 - c In the list, expand **CPU** and drag **Usage (%)** to the workspace on the right.
 - d From the threshold drop-down menu, select **Dynamic Threshold**.
 Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.
 - e In the **Symptom Definition Name** text box, type a name similar to **VM CPU Usage above trend**.
 - f From the criticality drop-down menu, select **Warning**.
 - g From the threshold drop-down menu, select **Above Threshold**.

- h Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

This Wait Cycle setting requires the symptom condition to be true for 3 collection cycles before the symptom is triggered. This wait avoids triggering the symptom when there is a short spike in CPU usage.

- i Click **Save**.

The dynamic symptom, which identifies when the usage is above the tracked trend, is added to the symptom list.

- 4 In the **Alert Definition Workspace** window, drag **VM CPU Usage above trend** from the symptom definition list to the symptom workspace on the right.

The Child-Virtual Machine symptom set is added to the symptom workspace.

- 5 In the symptoms set, configure the triggering condition so that when the symptom is true on half of the virtual machines in the group to which this alert definition is applied, the symptom set is true.
 - a From the value operator drop-down menu, select **>**.
 - b In the value text box, enter **50**.
 - c From the value type drop-down menu, select **Percent**.

Results

You defined the first symptom set for the alert definition.

What to do next

Add the host memory usage symptom to the alert definition. See [Add a Host Memory Usage Symptom to the Alert Definition](#).

Add a Host Memory Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add a second symptom to your vRealize Operations Manager alert definition after you add the first symptom. The second symptom is related to host memory usage for the hosts on which the accounting virtual machines operate.

Prerequisites

Add the virtual machine CPU usage symptom. See [Add a Virtual Machine CPU Usage Symptom to the Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Add Symptom Definitions**.

2 Configure the symptom related to host systems for the virtual machines.

- a From the **Defined On** drop-down menu, select **Self**.
- b From the **Symptom Definition Type** drop-down menu, select **Metric / Supermetric**.
- c Click the **Add** button to configure the new symptom.

3 Configure the host system symptom in the **Add Symptom Definition** workspace window.

- a From the **Base Object Type** drop-down menu, expand **vCenter Adapters** and select **Host System**.
- b In the metrics list, expand **Memory** and drag **Usage (%)** to the workspace on the right.
- c From the threshold drop-down menu, select **Dynamic Threshold**.

Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.

- d In the **Symptom Definition Name** text box, enter a name similar to **Host memory usage above trend**.
- e From the criticality drop-down menu, select **Warning**.
- f From the threshold drop-down menu, select **Above Threshold**.
- g Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

This Wait Cycle setting requires the symptom condition to be true for three collection cycles before the symptom is triggered. This wait avoids triggering the symptom when a short spike occurs in host memory usage.

- h Click **Save**.

The dynamic symptom identifies when the hosts on which the accounting virtual machines run are operating above the tracked trend for memory usage.

The dynamic symptom is added to the symptom list.

4 In the **Alert Definition Workspace** window, drag **Host memory usage above trend** from the symptoms list to the symptom workspace on the right.

The Self-Host System symptom set is added to the symptom workspace.

5 On the Self-Host System symptom set, from the value type drop-down menu for **This Symptom set is true when**, select **Any**.

With this configuration, when any of the hosts running accounting virtual machines exhibit memory usage that is above the analyzed trend, the symptom condition is true.

6 At the top of the symptom set list, from the **Match {operator} of the following symptoms** drop-down menu, select **Any**.

With this configuration, if either of the two symptom sets, virtual machine CPU usage or the host memory, are triggered, an alert is generated for the host.

Results

You defined the second symptom set for the alert definition and configured how the two symptom sets are evaluated to determine when the alert is generated.

What to do next

Add recommendations to your alert definition so that you and your engineers know how to resolve the alert when it is generated. See [Add Recommendations to the Alert Definition](#).

Add Recommendations to the Alert Definition

To resolve a generated alert for the accounting department's virtual machines, you provide recommendations so that you or other engineers have the information you need to resolve the alert before your users encounter performance problems.

As part of the alert definition, you add recommendations that include actions that you run from vRealize Operations Manager and instructions for making changes in vCenter Server that resolve the generated alert.

Prerequisites

Add symptoms to your alert definition. See [Add a Host Memory Usage Symptom to the Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, **Alert Impact**, and **Add Symptom Definitions**, click **Add Recommendations** and add the recommended actions and instructions.
- 2 Click **Add** and select an action recommendation to resolve the virtual machine alerts.
 - a In the **New Recommendation** text box, enter a description of the action similar to **Add CPUs to virtual machines**.
 - b From the **Actions** drop-down menu, select **Set CPU Count for VM**.
 - c Click **Save**.
- 3 Click **Add** and provide an instructive recommendation to resolve host memory problems similar to this example.

If this host is part of a DRS cluster, check the DRS settings to verify that the load balancing setting are configured correctly. If necessary, manually vMotion the virtual machines.

- 4 Click **Add** and provide an instructive recommendation to resolve host memory alerts.
 - a Enter a description of the recommendation similar to this example.
If this is a standalone host, add more memory to the host.
 - b To make the URL a hyperlink in the instructions, copy the URL, for example, <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html>, to your clipboard.
 - c Highlight the text in the text box and click **Create a hyperlink**.
 - d Paste the URL in the **Create a hyperlink** text box and click **OK**.
 - e Click **Save**.
- 5 In the **Alert Definition Workspace**, drag **Add CPUs to virtual machines, If this host is part of a DRS cluster**, and the **If this is a standalone host** recommendations from the list to the recommendation workspace in the order presented.
- 6 Click **Save**.

Results

You provided the recommended actions and instructions to resolve the alert when it is generated. One of the recommendations resolves the virtual machine CPU usage problem and the other resolves the host memory problem.

What to do next

Create a group of objects to use to manage your accounting objects. See [Create a Custom Accounting Department Group](#).

Create a Custom Accounting Department Group

To manage, monitor, and apply policies to the accounting objects as a group, you create a custom object group.

Prerequisites

Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).

Procedure

- 1 In the menu, click **Environment** and click the **Custom Groups** tab.
- 2 Click the **New Custom Group** icon to create a new custom group.
- 3 Type a name similar to **Accounting VMs and Hosts**.
- 4 From the **Group Type** drop-down menu, select **Department**.
- 5 From the **Policy** drop-down menu, select **Default Policy**.

When you create a policy, you apply the new policy to the accounting group.

- 6 In the Define membership criteria area, from the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Host System**, and configure the dynamic group criteria.

- a From the criteria drop-down menu, select **Relationship**.
- b From the relationships options drop-down menu, select **Parent of**.
- c From the operator drop-down menu, select **contains**.
- d In the **Object name** text box, enter **acct**.
- e From the navigation tree drop-down list, select **vSphere Hosts and Clusters**.

You created a dynamic group where host objects that are the host for virtual machines with acct in the virtual machine name are included in the group. If a virtual machine with acct in the object name is added or moved to a host, the host object is added to the group.

- 7 Click **Preview** in the lower-left corner of the workspace, and verify that the hosts on which your virtual machines that include acct in the object name appear in the **Preview Group** window.

- 8 Click **Close**.

- 9 Click **Add another criteria set**.

A new criteria set is added with the OR operator between the two criteria sets.

- 10 From the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Virtual Machine**, and configure the dynamic group criteria.

- a From the criteria drop-down menu, select **Properties**.
- b From the **Pick a property** drop-down menu, expand **Configuration** and double-click **Name**.
- c From the operator drop-down menu, select **contains**.
- d In the **Property value** text box, enter **acct**.

You created a dynamic group where virtual machine objects with acct in the object name are included in the group that depends on the presence of those virtual machines. If a virtual machine with acct in the name is added to your environment, it is added to the group.

- 11 Click **Preview** in the lower-left corner of the workspace, and verify that the virtual machines with acct in the object name are added to the list that also includes the host systems.

- 12 Click **Close**.

- 13 Click **OK**.

The Accounting VMs and Hosts group is added to the Groups list.

Results

You created a dynamic object group that changes as virtual machines with acct in their names are added, removed, and moved in your environment.

What to do next

Create a policy that determines how vRealize Operations Manager uses the alert definition to monitor your environment. See [Create a Policy for the Accounting Alert](#).

Create a Policy for the Accounting Alert

To configure how vRealize Operations Manager evaluates the accounting alert definition in your environment, you configure a policy that determines behavior so that you can apply the policy to an object group. The policy limits the application of the alert definition to only the members of the selected object group.

When an alert definition is created, it is added to the default policy and enabled, ensuring that any alert definitions that you create are active in your environment. This alert definition is intended to meet the needs of the accounting department, so you disable it in the default policy and create a new policy to govern how the alert definition is evaluated in your environment, including which accounting virtual machines and related hosts to monitor.

Prerequisites

- Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).
- Verify that you created a group of objects that you use to manage your accounting objects. See [Create a Custom Accounting Department Group](#).

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Policies**.
- 2 Click the **Policy Library** tab.
- 3 Click **Add New Policy**.
- 4 Type a name similar to **Accounting Objects Alerts Policy** and provide a useful description similar to the following example.

This policy is configured to generate alerts when Accounting VMs and Hosts group objects are above trended CPU or memory usage.

- 5 Select **Default Policy** from the **Start with** drop-down menu.
- 6 On the left, click **Customize Alert / Symptom Definitions** and disable all the alert definitions except the new Acct VM CPU early warning alert.
 - a In the Alert Definitions area, click **Actions** and select **Select All**.
The alerts on the current page are selected.
 - b Click **Actions** and select **Disable**.
The alerts indicate Disabled in the State column.

- c Repeat the process on each page of the alerts list.
- d Select **Acct VM CPU early warning** in the list, click **Actions** and select **Enable**.

The Acct VM CPU early warning alert is now enabled.

7 On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.

8 Click **Save**.

Results

You created a policy where the accounting alert definition exists in a custom policy that is applied only to the virtual machines and hosts for the accounting department.

What to do next

Create an email notification so that you learn about alerts even you when you are not actively monitoring vRealize Operations Manager. See [Configure Notifications for the Department Alert](#).

Configure Notifications for the Department Alert

To receive an email notification when the accounting alert is generated, rather than relying on your ability to generally monitor the accounting department objects in vRealize Operations Manager, you create notification rules.

Creating an email notification when accounting alerts are triggered is an optional process, but it provides you with the alert even when you are not currently working in vRealize Operations Manager.

Prerequisites

- Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).
- Verify that standard email outbound alerts are configured in your system. See [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1** In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.
- 2** Click **Notification Settings** and click the plus sign to add a notification rule.
- 3** Configure the communication options.
 - a In the **Name** text box, type a name similar to **Acct Dept VMs or Hosts Alerts**.
 - b From the **Select Plug-In Type** drop-down menu, select **StandardEmailPlugin**.
 - c From the **Select Instance** drop-down menu, select the standard email instance that is configured to send messages.

- d In the **Recipients** text box, type your email address and the addresses of other recipients responsible for the accounting department alerts. Use a semicolon between recipients.
- e Leave the **Notify again** text box blank.

If you do not provide a value, the email notice is sent only once. This alert is a Risk alert and is intended as an early warning rather than requiring an immediate response.

You configured the name of the notification when it is sent to you and the method that is used to send the message.

- 4 In the Filtering Criteria area, configure the accounting alert notification trigger.
 - a From the **Notification Trigger** drop-down menu, select **Alert Definition**.
 - b Click **Click to select Alert Definition**.
 - c Select **Acct VM CPU early warning** and click **Select**.
- 5 Click **Save**.

Results

You created a notification rule that sends you and your designated engineers an email message when this alert is generated for your accounting department alert definition.

What to do next

Create a dashboard with alert-related widgets so that you can monitor alerts for the accounting object group. See [Create a Dashboard to Monitor Department Objects](#).

Create a Dashboard to Monitor Department Objects

To monitor all the alerts related to the accounting department object group, you create a dashboard that includes the alert list and other widgets. The dashboard provides the alert data in a single location for all related objects.

Creating a dashboard to monitor the accounting virtual machines and related hosts is an optional process, but it provides you with a focused view of the accounting object group alerts and objects.

Prerequisites

Create an object group for the accounting department virtual machines and related objects. See [Create a Custom Accounting Department Group](#).

Procedure

- 1 In the menu, click **Dashboards > Actions > Create Dashboard**.
- 2 In the Dashboard Configuration definition area, type a tab name similar to **Accounting VMs and Hosts** and configure the layout options.
- 3 Click **Widget List** and drag the following widgets to the workspace.
 - **Alert List**

- **Efficiency**
- **Health**
- **Risk**
- **Top Alerts**
- **Alert Volume**

The blank widgets are added to the workspace. To change the order in which they appear, you can drag them to a different location in the workspace.

4 On the Alert List widget title bar, click **Edit Widget** and configure the settings.

- a In the **Title** text box, change the title to **Acct Dept Alert List**.
- b For the **Refresh Content** option, select **On**.
- c Type **Accounting** in the **Search** text box and click **Search**.

The Accounting value corresponds to the name of the object group for the accounting department virtual machines and related hosts.

- d In the filtered resource list, select the **Accounting VMs and Hosts** group.

The Accounting VMs and Hosts group is identified in the Selected Resource text box.

- e Click **OK**.

The Acct Dept Alert List is now configured to display alerts for the Accounting VMs and Hosts group objects.

5 Click **Widget Interactions** and configure the following interactions.

- a For Acct Dept Alert List, leave the selected resources blank.
- b For Top Alerts, Health, Risk, Efficiency, and Alert Volume select **Acct Dept Alert List** from the **Selected Resources** drop-down menu.
- c Click **Apply Interactions**.

With the widget interaction configured in this way, the select alert in the Acct Dept Alert List is the source for the data in the other widgets. When you select an alert in the alert list, the Health, Risk, and Efficiency widgets display alerts for that object, Top Alerts displays the topic issues affecting the health of the object, and Alert Volume displays an alert trend chart.

6 Click **Save**.

Results

You created a dashboard that displays the alerts related to the accounting virtual machines and hosts group, including the Risk alert you created.

Alerts Group

For easy and better management of alerts, you can arrange them as a group as per your requirement.


It is complicated to identify a problem in large environments as you receive different kind of alerts. To manage alerts easily, group them by their definitions.

For example, there are 1000 alerts in your system. To identify different types of alerts, group them based on their alert definitions. It is also easy to detect the alert having the highest severity in the group.

When you group alerts, you can see the number of times the alerts having the same alert definition are triggered. By grouping alerts, you can perform the following tasks easily and quickly:

- Find the noisiest alert: The alert that has triggered maximum number of times is known as the noisiest alert. Once you find it, you can disable it to avoid further noise.
- Filter alerts: You can filter alerts based on a substring in alert definitions. The result shows the group of alerts that contain the substring.

Note

- If you cancel or disable an alert group, the alerts are not canceled instantly. It might take some time if the group is large.
 - Only one group can be expanded at a time.
 - The number next to the group denotes the number of alerts in that particular group.
 - The criticality sign  indicates the highest level of severity of an alert in a group.
-

Grouping Alerts

You can group alerts by time, criticality, definition, and object type.

To group alerts:

Procedure

- 1 In the menu, click **Alerts**.
- 2 Select from the various options available from the **Group By** drop-down menu.

Disable Alerts

In an alerts group, you can disable an alert by a single click.

To disable an alert, in the menu, click **Alerts** and then in the left pane, click **All Alerts**. Select the alert name from the data grid, and click **Actions > Disable**.

The alerts can be disabled by two methods:

- **Disable Alert in All Policies:** You disable the alert for all the objects for all the policies.
- **Disable Alert in Selected Policies:** You disable the alert for the objects having the selected policy. Note that this method will work only for objects with alerts.

Configuring Actions

Actions are the ability to update objects or read data about objects in monitored systems, and are commonly provided in vRealize Operations Manager as part of a solution. The actions added by solutions are available from the object Actions menu, list and view menus, including some dashboard widgets, and can be added to alert definition recommendations.

The possible actions include read actions and update actions.

The read actions retrieve data from the target objects.

The update actions modifies the target objects. For example, you can configure an alert definition to notify you when a virtual machine is experiencing memory issues. Add an action in the recommendations that runs the Set Memory for Virtual Machine action. This action increases the memory and resolves the likely cause of the alert.

To see or use the actions for your vCenter Server objects, you must enable actions in the vCenter Adapter for each monitored vCenter Server instance. Actions can only be viewed and accessed if you have the required permissions.

List of vRealize Operations Manager Actions

The list of actions includes the name of the action, the objects that each one modifies, and the object levels at which you can run the action. You use this information to ensure that you correctly apply the actions as alert recommendations and when the actions are available in the **Actions** menu.

Actions and Modified Objects

vRealize Operations Manager actions make changes to objects in your managed vCenter Server instances.

When you grant a user access to actions in vRealize Operations Manager, that user can take the granted action on any object that vRealize Operations Manager manages.

Action Object Levels

The actions are available when you work with different object levels, but they modify only the specified object. If you are working at the cluster level and select **Power On VM**, all the virtual machines in the cluster for which you have access permission are available for you to run the action. If you are working at the virtual machine level, only the selected virtual machine is available.

Table 2-3. vRealize Operations Manager Actions Affected Objects

Action	Modified Object	Object Levels
Rebalance Container	Virtual Machines	<ul style="list-style-type: none"> ■ Data Center ■ Custom Data Center
Delete Idle VM	Virtual Machines	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set DRS Automation	Cluster	<ul style="list-style-type: none"> ■ Clusters
Move VM	Virtual Machine	<ul style="list-style-type: none"> ■ Virtual Machines
Power Off VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Shut Down Guest OS for VM	Virtual Machine VMware Tools must be installed and running on the target virtual machines to run this action.	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Power On VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Delete Powered Off VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set Memory for VM and Set Memory for VM Power Off Allowed	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set Memory Resources for VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set CPU Count for VM and Set CPU Count for VM Power Off Allowed	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set CPU Resources for VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set CPU Count and Memory for VM and Set CPU Count and Memory for VM Power Off Allowed	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines

Table 2-3. vRealize Operations Manager Actions Affected Objects (continued)

Action	Modified Object	Object Levels
Delete Unused Snapshots for VM	Snapshot	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Delete Unused Snapshots for Datastore	Snapshot	<ul style="list-style-type: none"> ■ Clusters ■ Datastores ■ Host Systems

Actions Supported for Automation

Recommendations can identify ways to remediate problems indicated by an alert. Some of these remediations can be associated with actions defined in your vRealize Operations Manager instance. You can automate several of these remediation actions for an alert when that recommendation is the first priority for that alert.

You enable actionable alerts in your policies. By default, automation is disabled in policies. To configure automation for your policy, in the menu, click **Administration > Policies > Policy Library**. Then, you edit a policy, access the **Alert / Symptom Definitions** workspace, and select **Local** for the **Automate** setting in the Alert / Symptom Definitions pane.

When an action is automated, you can use the **Automated** and **Alert** columns in **Administration > History > Recent Tasks** to identify the automated action and view the results of the action.

- vRealize Operations Manager uses the **automationAdmin** user account to trigger automated actions. For these automated actions that are triggered by alerts, the Submitted By column displays the **automationAdmin** user.
- The Alert column displays the alert that triggered the action. When an alert is triggered that is associated to the recommendation, it triggers the action without any user intervention.

The following actions are supported for automation:

- Delete Powered Off VM
- Delete Idle VM
- Move VM
- Power Off VM
- Power On VM
- Set CPU Count And Memory for VM
- Set CPU Count And Memory for VM Power Off Allowed
- Set CPU Count for VM
- Set CPU Count for VM Power Off Allowed
- Set CPU Resources for VM

- Set Memory for VM
- Set Memory for VM Power Off Allowed
- Set Memory Resources for VM
- Shut Down Guest OS for VM

Roles Needed to Automate Actions

To automate actions, your role must have the following permissions:

- Create, edit, and import policies in **Administration > Policies > Policy Library**.
- Create, clone, edit, and import alert definitions in **Alerts > Alert Settings > Alert Definitions**.
- Create, edit, and import recommendation definitions in **Alerts > Alert Settings > Recommendations**.

Important You set the permissions used to run the actions separately from the alert and recommendation definition. Anyone who can modify alerts, recommendations, and policies can also automate the action, even if they do not have permission to run the action.

For example, if you do not have access to the Power Off VM action, but you can create and modify alerts and recommendations, you can see the Power Off VM action and assign it to an alert recommendation. Then, if you automate the action in your policy, vRealize Operations Manager uses the `automationAdmin` user to run the action.

Example Action Supported for Automation

For the Alert Definition named `Virtual machine has chronic high CPU workload leading to CPU stress`, you can automate the action named `Set CPU Count for VM`.

When CPU stress on your virtual machines exceeds a critical, immediate, or warning level, the alert triggers the recommended action without user intervention.

Integration of Actions with vRealize Automation

vRealize Operations Manager restricts actions on objects that vRealize Automation manages, so that the actions do not violate any constraints set forth by vRealize Automation.

When objects in your environment are managed by vRealize Automation, actions in vRealize Operations Manager are not available on those objects. For example, if a host or parent object is being managed by vRealize Automation, actions are not available on that object.

This behavior is true for all actions, including **Power Off VM**, **Move VM**, **Rebalance Container**, and so on.

You cannot turn on or turn off the exclusion of actions on vRealize Automation managed objects.

Actions Determine Whether Objects Are Managed

Actions check the objects in the vRealize Automation managed resource container to determine which objects are being managed by vRealize Automation.

- Actions such as Rebalance Container check the child objects of the data center container or custom data center container to determine whether the objects are managed by vRealize Automation. If the objects are being managed, the action does not appear on those objects.
- The Move VM action checks whether the virtual machine to be moved is being managed by vRealize Automation.

Is the Virtual Machine Managed?	Result of Move VM Action
Yes	The Move VM action does not appear in the vRealize Operations Manager user interface for that virtual machine.
No	The Move VM action moves the virtual machine to a new host, datastore, or new host and datastore. The Move VM action does not check whether the new host or datastore is being managed by vRealize Automation.

- The Delete Snapshots action checks whether the virtual machine or datastore is being managed by vRealize Automation.

Actions on Objects that vRealize Automation Does Not Manage

For a host or parent object that is not managed by vRealize Automation, only the virtual machines that are not being managed by vRealize Automation appear in the action dialog, and you can only take action on the virtual machines that are not being managed by vRealize Automation. If all child objects are being managed by vRealize Automation, the user interface displays the message *No objects are eligible for the selected action.*

If You Attempt to Run an Action on Multiple Objects

If you select multiple objects and attempt to run an action, such as Power Off VM, only the objects that are not being managed by vRealize Automation, which might include a subset of the virtual machines, appear in the Power Off VM action dialog box.

Working with Actions That Use Power Off Allowed

Some of the actions provided with vRealize Operations Manager require the virtual machines to shut down or power off, depending on the configuration of the target machines, to run the actions. You should understand the impact of the Power Off Allowed option before running the actions so that you select the best options for your target virtual machines.

Power Off and Shut Down

The actions that you can run on your vCenter Server instances include actions that shut down virtual machines and actions that power off virtual machines. It also includes actions where the virtual machine must be in a powered off state to complete the action. Whether the VM is shut down or powered off depends on how it is configured and what options you select when you run the action.

The shut-down action shuts down the guest operating system and then powers off the virtual machine. To shut down a virtual machine from vRealize Operations Manager, the VMware Tools must be installed and running on the target objects.

The power off action turns off the VM without regard for the state of the guest operating system. In this case, if the VM is running applications, your user might lose data. After the action is finished, for example, modifying the CPU count, the virtual machine is returned to the power state it was in when the action began.

Power Off Allowed and VMware Tools

For the actions where you are increasing the CPU count or the amount of memory on a VM, some operating systems support the actions if the Hot Plug is configured on the VM. For other operating systems, the virtual machine must be in a powered off state to change the configuration. To accommodate this need where the VMware Tools is not running, the Set CPU Count, Set Memory, and Set CPU Count and Memory actions include the Power Off Allowed option.

If you select Power Off Allowed, and the machine is running, the action verifies whether VMware Tools is installed and running.

- If VMware Tools is installed and running, the virtual machine is shut down before completing the action.
- If VMware Tools is not running or not installed, the virtual machine is powered off without regard for the state of the operating system.

If you do not select Power Off Allowed and you are decreasing the CPU count or memory, or the hot plug is not enabled for increasing the CPU count or memory, the action does not run and the failure is reported in Recent Tasks.

Power Off Allowed When Changing CPU Count or Memory

When you run the actions that change the CPU count and the amount of memory, you must consider several factors to determine if you want to use the Power Off Allowed option. These factors include whether you are increasing or decreasing the CPU or memory and whether the target virtual machines are powered on. If you increase the CPU or memory values, whether hot plug is enabled also affects how you apply the option when you run the action.

How you use Power Off Allowed when you are decreasing the CPU count or the amount of memory depends on the power state of the target virtual machines.

Table 2-4. Decreasing CPU Count and Memory Behavior Based On Options

Virtual Machine Power State	Power Off Allowed Selected	Results
On	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, decreases the CPU or memory, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, decreases the CPU or memory, and powers the machine back on.</p>
On	No	The action does not run on the virtual machine.
Off	Not applicable. The virtual machine is powered off.	The action decreases the value and leaves the virtual machine in a powered off state.

How you use Power Off Allowed when you are increasing the CPU count or the amount of memory depends on several factors, including the state of the target virtual machine and whether hot plug is enabled. Use the following information to determine which scenario applies to your target objects.

If you are increasing the CPU count, you must consider the power state of the virtual machine and whether CPU Hot Plug is enabled when determining whether to apply Power Off Allowed.

Table 2-5. Increasing CPU Count Behavior.

Virtual Machine Power State	CPU Hot Plug Enabled	Power Off Allowed Selected	Results
On	Yes	No	The action increases the CPU count to the specified amount.
On	No	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, increases the CPU count, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, increases the CPU count, and powers the machine back on.</p>
Off	Not applicable. The virtual machine is powered off.	Not required.	The action increases the CPU count to the specified amount.

If you are increasing the memory, you must consider the power state of the virtual machine, whether Memory Hot Plug is enabled, and whether there is a Hot Memory Limit when determining how to apply Power Off Allowed.

Table 2-6. Increasing Memory Amount Behavior

Virtual Machine Power State	Memory Hot Plug Enabled	Hot Memory Limit	Power Off Allowed Selected	Results
On	Yes	New memory value \leq hot memory limit	No	The action increases the memory the specified amount.
On	Yes	New memory value $>$ hot memory limit	Yes	If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on. If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on.
On	No	Not applicable. The hot plug is not enabled.	Yes	If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on. If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on.
Off	Not applicable. The virtual machine is powered off.	Not applicable.	Not required	The action increases the memory the specified amount.

Configuring and Using Workload Optimization

3

Workload Optimization provides for moving virtual compute resources and their file systems dynamically across datastore clusters within a data center or custom data center.

Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.

Workload Optimization further enables you potentially to automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention automatically runs an action, a data center performs at optimum.

vRealize Automation Integration

When you add an instance to a vRealize Automation adapter or solution pack as well as to a vCenter Server adapter instance that is connected to the vRealize Automation server, using vRealize Automation-managed resources, vRealize Operations Manager automatically adds a custom data center for the vCenter Server, using vRealize Automation-managed resources.

On the vRealize Operations Manager side, to get the day2 chain configured, you must make the following initial configurations:

- 1 In vCenter Server, **Administration -> Solutions** and then add the VMware vSphere adapter instance for the vCenter Server that is configured as an endpoint in vRealize Automation Server.
- 2 In vCenter Server, **Administration -> Solutions** and then add the VMware vRealize Automation adapter instance for the server that will appear in the vRealize Operations Manager and vRealize Automation integration day2 chain.

vRealize Operations Manager can manage workload placement and optimization for the custom data centers that reside in vRealize Automation-managed clusters.

However, vRealize Operations Manager is not permitted to set tag policies for the custom data center. (At the Workload Optimization screen, the Business Intent window is not operational for vRealize Automation custom data centers.) When rebalancing a vRealize Automation custom data center, vRealize Operations Manager uses all applicable policies and placement principles from both systems: vRealize Automation and vRealize Operations Manager. For more information on configuring vRealize Automation to work with vRealize Operations Manager, see [vRealize Automation Solution](#). For complete information on creating and managing vRealize Automation custom data centers that are managed by vRealize Operations Manager, see the vRealize Automation documentation.

This chapter includes the following topics:

- [Configuring Workload Optimization](#)
- [Using Workload Optimization](#)

Configuring Workload Optimization

Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

- 1 Configure the Workload Automation Details. See [Workload Automation Details](#).
- 2 Tag VMs for cluster placement. See [Business Intent - Host-Based Virtual Machine Placement](#) and [Business Intent: Tag-Based VM Placement in Clusters](#).
- 3 If you do not use the AUTOMATE function in the Optimization Recommendation pane at the Workload Automation screen, configure the two Workload Optimization alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the actions calculated by Workload Optimization are run automatically. See [Configuring Workload Optimization Alerts](#)

Prerequisites

Workload Optimization acts on objects associated with the VMware vSphere Solution that connects vRealize Operations Manager to one or more vCenter Server instances. The virtual objects in this environment include a vCenter Server, data centers and custom data centers, cluster compute and storage resources, host systems, and virtual machines. Specific requirements:

- A vCenter Adapter configured with the actions enabled for each vCenter Server instance.
- A vCenter Server instance with at least two datastore clusters with sDRS enabled and fully automated.
- Any non-datastore clusters must have DRS enabled and fully automated
- Storage vMotion must be set to ON at Workload Automation Details. The default is On.
- You must have permission to access all objects in the environment.

Design Considerations

The following rules constrain the possible computer and storage resource moves that can be performed.

Note When vRealize Operations Manager suggests that you optimize clusters in a data center, the system does not guarantee it can run an optimization action. vRealize Operations Manager analytics can determine that optimization is desirable and can create a rebalancing plan. However, the system cannot automatically identify all the architectural constraints that may be present. Such constraints may prevent an optimization action, or cause an action in progress to fail.

- Moving compute and storage resources is allowed only within, not across data centers or custom data centers.
- Storage resources cannot be moved across non-datastore clusters. Storage can move only across datastore clusters that have sDRS fully automated.
- Compute-resource-only moves are permitted through shared storage.
- Virtual machines defined with affinity rules or anti-affinity rules are not to be moved.
- Virtual machines cannot be moved when residing on a local datastore, unless a storage swap exists on the local datastore.
- Virtual machines cannot be moved if they have data residing across multiple datastore clusters. Compute-only moves with similar shared storage are not permitted.
- A virtual machine cannot have data that resides across different storage types. For example, if a virtual machine has a VM disk on a datastore and a second VM disk on a datastore cluster, the virtual machine does not move, even when the datastore is shared with the destination or has swap on it.
- A virtual machine can use RDM so long as the destination datastore cluster can access the RDM LUN.
- A virtual machine can implement VM disks on multiple datastores inside a single datastore cluster.
- Workload Optimization may suggest moving virtual machines that are protected by vSphere Replication or Array Based Replication. You must ensure that all the clusters within a selected data center or custom data center have replication available. You can set up DRS affinity rules on virtual machines that you do not want moving across clusters.

Business Intent: Tag-Based VM Placement in Clusters

You can use vCenter Server tagging to tag VMs and associated clusters, respectively, with specific tags. These tags define - for a given cluster - the set of VMs that is placed with that cluster and remains within the cluster. When the system runs an optimization action, it uses VM-to-cluster tag matching to ensure that VMs are moved to - or stay with - the appropriate cluster.

To edit Business Intent values, you must have privileges for Administration -> Configuration -> Workload Placement Settings -> Edit.

Using Tags for Cluster Flexibility

When configuring custom data centers and clusters without tags, you configure CDCs as relatively homogeneous. All cluster resources must support, for example, the same OS or the same security requirements so that optimization actions do not place VMs in the wrong cluster.

The tagging approach enables you to define zones of infrastructure within cluster boundaries. For example, you can ensure that during workload optimization actions, Windows VMs are moved only to Windows-licensed clusters and Oracle VMs are moved only to Oracle-licensed clusters. Similarly, you can enable tiers of service in an application, where "Tier 1" VMs are moved only to Tier 1 clusters running business-critical applications. Other examples include separating VMs according to OS, or creating network boundaries.

VMs and clusters can be tagged with more than one tag. VMs with multiple tags are placed only on clusters with all matching tags.

Note VM-to-cluster tagging is not the same as host-based VM tagging. See [Business Intent - Host-Based Virtual Machine Placement](#).

vCenter Server tags are implemented as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name.

Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

The system provides several preset categories at the Business Intent Workspace:

- Operating System
- Environment
- Tier
- Network
- Other

These categories represent potential business intent in gathering VMs into various associations. You are free to remove a category or add a new one that works for your environment.

Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

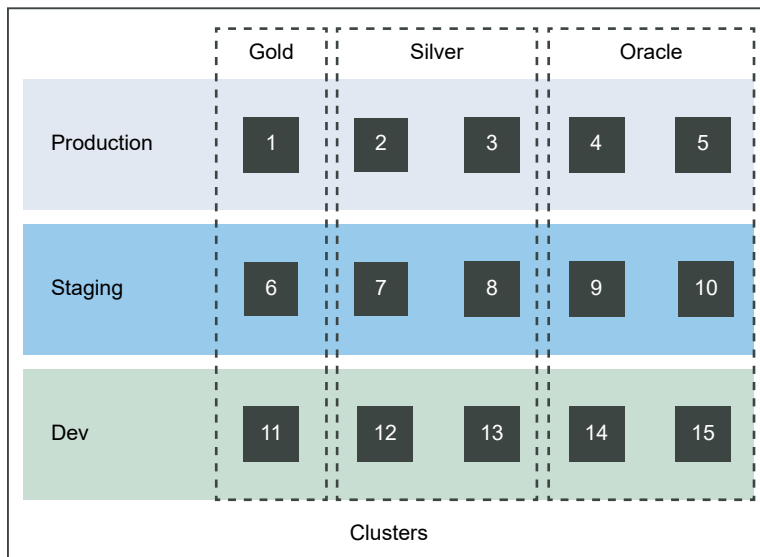
In vRealize Operations Manager, you assign category and name tags in Policies, at the Business Intent workspace.

Tagging Considerations

- You can choose either cluster-tag-based placement or host-based placement in the same data center or custom data center, but not both. If you select cluster-tag-based placement, host tags are ignored. Conversely, if you choose host-tag-based placement, cluster tags are ignored.
- If a VM is tagless, the system attempts to move it to a tagless cluster.

Tag Implementation Example: Cluster Zones of Service and Licensing

The following example shows how an administrator assigned tags to clusters and VMs to create zones within a data center:



Data Center A

Using vCenter Server, the administrator sets up these tag categories and associated tag names:

- Environment: Production, Staging, Dev
- Service Tier: Gold, Silver
- Licensing: Oracle

Data Center A includes 15 clusters. The administrator tags the clusters and VMs in those clusters as follows:

Cluster	Environment	Service Tier	Licensing
1	Production	Gold	
2, 3	Production	Silver	
4, 5	Production		Oracle
6	Staging	Gold	
7, 8	Staging	Silver	

Cluster	Environment	Service Tier	Licensing
9, 10	Staging		Oracle
11	Dev	Gold	
12, 13	Dev	Silver	
14, 15	Dev		Oracle

Opening the vRealize Operations Manager policies to Tag-Based VM Placement in the Business Intent window, the administrator prioritizes the Environment: Production and Service Tier: Gold category-tag combinations. Because the Optimization policies emphasize balance, clusters with those tags are balanced first.

Business Intent - Host-Based Virtual Machine Placement

Use host-based VM placement to tie your VMs more closely to your infrastructure. By using vCenter Server to tag hosts and VMs with specific tags, you make certain that when the system runs an optimization, it uses VM-to-host tag matching to ensure that VMs are moved to - or stay with - the appropriate host.

Using Tags to Enhance Structure

When configuring data centers or custom data centers without tags, you configure clusters and their hosts as relatively homogenous. All cluster resources must support, for example, the same OS or the same security requirements so that optimization actions do not place VMs in the wrong cluster.

The tagging approach enables you to define zones of infrastructure within cluster boundaries. VM-to-cluster tagging, for example, allows you to tag VMs and clusters to assure that Windows VMs are moved only to Windows-licensed clusters and Oracle VMs are moved only to Oracle-licensed clusters.

With host-based VM placement (VM-to-host tagging), you bind your VMs to individual hosts rather than clusters.

vCenter Server tags are implemented as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. You can define many keys and values in vCenter Server, but choose a subset to be considered in the Business Intent pane of the Workload Optimization screen (**Home -> Optimize Performance -> Workload Optimization**).

Note If you choose host-based placement in the Business Intent pane, the system - after getting confirmation from you - disables conflicting user-created affinity rules. Then, as you define host-VM tagging relationships in the Business Intent pane, vRealize Operations Manager automatically creates the required affinity rules, saving you the manual effort. So, for example, suppose you configure a tag in the Business Intent pane that requires VM1 to remain with Host1. If there exists a user-configured affinity rule keeping VM1 with Host2, the system disables the rule. However, if another user-configured affinity rule dictates that VM2 remains with Host2, the system does not change that rule.

Additional Considerations

- You are not permitted to employ both VM-to-cluster tagging and VM-to-host tagging in the same data center or custom data center - only one tagging method or the other. If you select host-based VM placement, any cluster tags are ignored.
- With host-based VM placement, only one category and one tag per VM is allowed per VM.
- A tagless VM can be sent to any host, even a tagged host.
- A host with multiple tags is treated as tagless.
- Even if all workloads are balanced, if there is also a tag violation, the system is by definition not optimized.
- The system does not consider any tags of storage - that is, datastores or datastore clusters.

Business Intent Workspace

You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. vRealize Operations Manager can be configured to leverage tags to define business-related placement constraints: VMs can only be placed on hosts/clusters with matching tags.

Where You Find Business Intent

From the Home page, click the chevron next to Optimize Performance on the left. Click Workload Optimization, select a data center or custom data center from the top row, and click **Edit** in the Business Intent window.

To edit Business Intent values, you must have privileges for Administration -> Configuration -> Workload Placement Settings -> Edit.

Establishing Business Intent

Tags are implemented in vCenter Server as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

To specify tags considered for placement, first select the radio button for the type of object you want to associate with VMs in this business intent session: Clusters or Hosts.

The system provides several suggested categories. These categories are only suggestions. You must specify the actual categories in vCenter Server after you expand the section for a suggested category . For example, in section "Tier", you can specify the actual vCenter Server tag category that represents tier semantics, for instance, "service level".

- Operating System
- Environment
- Tier

- Network
- Other

Any actual categories you specify must first be created in vCenter Server.

Then you can associate tagged VMs with clusters or hosts, based on the rules for each type of tagging.

- 1 Click the chevron to the left of the first suggested category. A **tag category** field appears.
- 2 Click the drop-down menu indicator and choose a category from the list defined in vCenter Server.
- 3 Click the drop-down menu indicator in the Tag Name (Optional) field and choose a tag name from the list defined in vCenter Server.
- 4 Click **Include Tag**. All VMs with that tag are associated with the category.

Rules for Host-Based Placement

To set host level placement constraints, vRealize Operations Manager automatically creates and manages DRS rules. All conflicting user-created DRS rules are DISABLED.

These rules include the following:

- Any VM-VM affinity and anti-affinity rules.
- Any VM-Host affinity and anti-affinity rules.

You must check the selection box next to the statement, "I understand that vRealize Operations will disable all my current and future DRS rules".

See also [Business Intent - Host-Based Virtual Machine Placement](#) .

Rules for Cluster-Based Placement

See [Business Intent: Tag-Based VM Placement in Clusters](#).

Configuring Workload Optimization Alerts

vRealize Operations Manager provides two preconfigured alerts designed to work with the Workload Optimization feature. You must take additional action in the Policies area to turn on the alerts and automate them so that predetermined actions are run when the alerts fire.

The following preconfigured alerts are designed to work with the Workload Optimization feature:

- Data center performance can potentially be optimized in one or more clusters.
- Custom data center performance can potentially be optimized in one or more clusters.

The preconfigured alerts fire only if the AUTOMATE function is not turned on at the Workload Optimization screen. (**Home -> Optimize Performance -> Workload Optimization**).

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI pages and manage vCenter Server objects.

Procedure

- 1 Select **Administration** from the menu, then **Policies** from the left pane.
- 2 Click **Policy Library** and select the policy that includes settings for the relevant data centers and custom data centers, for example, **vSphere Solution's Default Policy**.
- 3 Click **Edit**.
- 4 Click #6 on the lower left, Alert/Symptom Definitions.
- 5 Search on "can potentially be optimized" to locate the two alerts you want.
- 6 The alerts are turned ON by default/inheritance (see the State column).
- 7 The alerts are not automated by default/inheritance (see the Automate column). To automate the alerts, click the menu symbol to the right of the inherited value and select the green check mark.

Results

Workload Optimization is fully automated for your environment.

What to do next

To confirm that actions are taken automatically, monitor rebalance activity at the Workload Optimization screen.

Using Workload Optimization

Use the Workload Optimization UI pages to monitor optimizing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and run actions directly.

vRealize Operations Manager monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use optimization functions to distribute a workload differently in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects of interest.

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see .

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see the *vRealize Operations Manager User Guide*.

The following examples demonstrate the primary ways you can use Workload Optimization to keep your data centers balanced and performing their best.

Example: Run Workload Optimization

As a virtual infrastructure administrator or other IT professional, you use Workload Optimization functions to identify points of resource contention or imbalance. In this example, you manually run an optimization action to consolidate demand.

When you log into vRealize Operations Manager, you see the Quick Start page. In the left-most column, Optimize Performance, is the alert 3 DATA CENTERS REQUIRING OPTIMIZATION.

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

Procedure

- 1 Click **Workload Optimization** in the Optimize Performance column.

The Workload Optimization page appears. Data centers are grouped by Criticality, with the three troubled data centers appearing in a carousel across the top of the page: DC-Bangalore-18, DC-Bangalore-19, DC-Bangalore-20. A Not Optimized badge appears in the lower right corner of each graphic.

- 2 If no data center is preselected, select DC-Bangalore-18 from the carousel.

Comprehensive data about the state of the data center follows.

- 3 Based on the available data, you determine an optimization action is required.

CPU workloads can be consolidated such that a host in Cluster 3 can be freed up.

Table 3-1. Panes and Widgets

Pane	Contents
Workload Optimization	Status shows as Not Optimized. A system message says, "You can consolidate workloads to maximize usage and potentially free up 1 host." The message reflects that you have set policies to emphasize consolidation as a goal in optimization moves. The system is saying you can free up a host through consolidation.
Settings	The current policy is Consolidate. The system advises: Avoid Performance Issues, Consolidate Workloads.
Cluster Workloads	Cluster 1 CPU Workload is 16%. Cluster 2 CPU Workload is 29%. Cluster 3 CPU Workload is 14%. Cluster 4 CPU Workload is 22%.

- 4 Click **OPTIMIZE NOW** in the Workload Optimization pane.

The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.

- 5 If you are satisfied with the projected results of the optimization action, click **NEXT**.

The dialog box updates to show the planned moves.

- 6 Review the optimization moves, then click **BEGIN ACTION**.

The system runs the compute and storage resource moves.

Results

The optimization action moved compute and storage resources from some clusters to other clusters in the data center, and so freed up a host on one cluster.

Note The Workload Optimization page refreshes every five minutes. Depending on when you run an optimization action, the system might not reflect the result for up to five minutes, or longer when longer-running actions extend the processing time.

What to do next

To confirm that your optimization action was completed, go to the Recent Tasks page by selecting **Administration** on the top menu, and clicking **History > Recent Tasks** in the left pane. In the Recent Tasks page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

Note Sometimes an optimizing action may be suggested, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that suggested optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation might incur stress in the future, then consolidation is not suggested.

Example: Schedule a Repeating Optimization Action

As a virtual infrastructure administrator or other IT professional, you determine that compute and storage resources in a given data center are volatile and a regularly scheduled optimization action can address the problem.

vRealize Operations Manager monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization page. Depending on what appears, you may determine that you must schedule optimization functions to distribute a workload more evenly in a data center or custom data center.

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

Procedure

- 1 From the Home screen, click **Optimize Performance > Workload Optimization** in the left pane.
- 2 From the carousel of data centers across the top of the page, select a data center for which you want to schedule repeated optimization actions.
- 3 In the Workload Optimization pane, click **SCHEDULE**.
- 4 Give the schedule a name and choose a time zone.
- 5 Determine how often you want to repeat the optimization action and click the relevant **radio button** under Recurrence.

Depending on your selection under Recurrence, additional options appear to the right. In this instance, you choose to repeat the optimization daily.
- 6 Leave the current date and time.
- 7 Select the **Repeat every day** radio button.
- 8 Select the **Expire after** radio button and tick the counter up to 6.
- 9 Click **Save**.

Results

The optimization action repeats for six days, then stops.

At the Workload Optimization page, the Scheduled button appears in the upper right of the Workload Optimization pane if optimization actions are scheduled for the selected data center. If you want to edit or delete a schedule, click the **Scheduled** button. The Optimization Schedules page appears, where you can perform those actions.

Note If you schedule a number of optimization actions close together, and the optimization plans of two or more actions include overlapping functions, that is, they impact the same set of resources, the system shifts the actions into a queue. As a result, some actions may complete later than expected, with longer running actions and other potential system constraints extending the lag time. Optimization actions that do not overlap can run concurrently.

What to do next

To confirm that your optimization action was finished, go to the Recent Tasks screen by selecting **Administration** on the top menu, and clicking **History > Recent Task** in the left pane. In the Recent Tasks screen, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, filter on Event Source and enter the name of the scheduled optimization plan.

Note Because real-time data center resource contention is dynamic, the system calculates a new optimization plan each time the scheduled optimization action starts, but before it runs. The system does not run the action if the system determines that the data center container is balanced at this moment. On the Recent Tasks page, the name of the affected data center appears in the Object Name column, and the Message “The optimization of the selected container cannot be improved” appears under Details. Another possibility is that a scheduled optimization plan is attempted, but does not go forward. In this event - which is not the same as a "failed" action - the name of the affected data center also appears in the Object Name column.

Example: Run Workload Optimization from Recommended Actions

From the Home screen, click **Recommendations** under Optimize Performance - first column on the left. The Recommended Actions screen appears, with data center and custom data center errors highlighted. If a suggested optimization action is available, it appears in the bottom third of the screen, with details.

To run the action, click the blue **Run Action** arrow.

The screenshot displays the 'Recommended Actions' interface. On the left, a sidebar lists navigation options: Recommended Actions, Operations Overview, Capacity Overview, Workload Balance, Log Insight, and Business Management. The main panel is titled 'Recommended Actions' and includes a 'Select Object Type' dropdown menu with options: All vCenters, vCenter Server (1), Datacenter (1), Custom Datacenter, and Cluster Compute Resou. The 'Datacenter (1)' option is selected. Below the dropdown, a 'Health Status' section shows a red circle with the number '1' and the text 'Objects', indicating 1 Critical and 0 Immediate issues. To the right, a 'Worst Health' table lists the object 'CMBU_ESO_VC09_DC' with 1 alert. At the bottom, a 'Suggested Fix' section shows an action to 'Optimize the cluster by spreading the workload' for the same object, with a description 'Datacenter performance can potentially be optim...' and an alert type 'Virtualization/Hyperv... Alerts'.

Prerequisites

Ensure that you have all required permissions for accessing the Workload Optimization UI and managing vCenter Server objects.

Results

The system runs the proposed rebalancing action.

What to do next

The Workload Optimization screen appears, where you can review the results of the rebalancing actions. Additional information is available at the Recent Tasks page: in the menu, select **Administration**, then click **History > Recent Tasks** in the left pane. Choose the **Event Source** filter and enter part of the alert name, then search. If the action succeeded, the Event Source column shows Alert: *<alert name>*.

Configuring Policies

4

To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

This chapter includes the following topics:

- [Policies](#)
- [Operational Policies](#)
- [Types of Policies](#)
- [Using the Monitoring Policy Workspace to Create and Modify Operational Policies](#)

Policies

A policy is a set of rules that you define for vRealize Operations Manager to use to analyze and display information about the objects in your environment. You can create, modify, and administer policies to determine how vRealize Operations Manager displays data in dashboards, views, and reports.

How Policies Relate to Your Environment

vRealize Operations Manager policies support the operational decisions established for your IT infrastructure and business units. With policies, you control what data vRealize Operations Manager collects and reports on for specific objects in your environment. Each policy can inherit settings from other policies, and you can customize and override various analysis settings, alert definitions, and symptom definitions for specific object types, to support the service Level agreements and business priorities established for your environment.

When you manage policies, you must understand the operational priorities for your environment, and the tolerances for alerts and symptoms to meet the requirements for your business critical applications. Then, you can configure the policies so that you apply the correct policy and threshold settings for your production and test environments.

Policies define the settings that vRealize Operations Manager applies to your objects when it collects data from your environment. vRealize Operations Manager applies policies to newly discovered objects, such as the objects in an object group. For example, you have an existing VMware adapter instance, and you apply a specific policy to the group named World. When a user adds a new virtual machine to the vCenter Server instance, the VMware adapter reports the virtual machine object to vRealize Operations Manager. The VMware adapter applies the same policy to that object, because it is a member of the World object group.

To implement capacity policy settings, you must understand the requirements and tolerances for your environment, such as CPU use. Then, you can configure your object groups and policies according to your environment.

- For a production environment policy, a good practice is to configure higher performance settings, and to account for peak use times.
- For a test environment policy, a good practice is to configure higher utilization settings.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Table 4-1. Configurable Policy Rule Elements

Policy Rule Elements	Thresholds, Settings, Definitions
Workload	Configure symptom thresholds for Workload.
Time Remaining	Configure thresholds for the Time Remaining.
Capacity Remaining	Configure thresholds for the Capacity Remaining.
Maintenance Schedule	Sets a time to perform maintenance tasks.
Attributes	An attribute is a collectible data component. You can enable or disable metric, property, and super metric attributes for collection, and set attributes as key performance indicators (KPIs). A KPI is the designation of an attribute that indicates that the attribute is important in your own environment.
Alert Definitions	Enable or disable combinations of symptoms and recommendations to identify a condition that classifies as a problem.
Symptom Definitions	Enable or disable test conditions on properties, metrics, or events.

Privileges to Create, Modify, and Prioritize Policies

You must have privileges to access specific features in the vRealize Operations Manager user interface. The roles associated with your user account determine the features you can access and the actions you can perform.

To set the policy priority, on the Active Policies tab, click the policy row and drag it to place it at the desired priority in the list. The priority for the Default Policy is always designated with the letter D.

How Upgrades Affect Your Policies

After you upgrade vRealize Operations Manager from a previous version, you might find newly added or updated default settings of policies such as, new alerts and symptoms. Hence, you must analyze the settings and modify these settings to optimize them for your current environment. If you apply the policies used with a previous version of vRealize Operations Manager, the manually modified policy settings remain unaltered.

Policy Decisions and Objectives

Implementing policy decisions in vRealize Operations Manager is typically the responsibility of the Infrastructure Administrator or the Virtual Infrastructure Administrator, but users who have privileges can also create and modify policies.

You must be aware of the policies established to analyze and monitor the resources in your IT infrastructure.

- If you are a Network Operations engineer, you must understand how policies affect the data that vRealize Operations Manager reports on objects, and which policies assigned to objects report alerts and issues.
- If you are the person whose role is to recommend an initial setup for policies, you typically edit and configure the policies in vRealize Operations Manager.
- If your primary role is to assess problems that occur in your environment, but you do not have the responsibility to change the policies, you must still understand how the policies applied to objects affect the data that appears in vRealize Operations Manager. For example, you might need to know which policies apply to objects that are associated with particular alerts.
- If you are a typical application user who receives reports from vRealize Operations Manager, you must have a high-level understanding of the operational policies so that you can understand the reported data values.

Active Policies Tab for Policies

The **Active Policies** tab displays the policies associated with groups of objects. You can manage the active policies for the objects in your environment so that you can have vRealize Operations Manager analyze and display specific data about those objects in dashboards, views, and reports.

How the Active Policies Tab Works

Use the **Active Policies** tab to associate a policy with one or more object groups, and to set the default policy. You can view the locally defined settings for a policy, and the complete list of settings, which includes those that are inherited from the base policies that you select in the Add or Edit Policy workspace. You can assign any policy to be the default policy.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

To display the details for a selected policy, click the split bar to expand the pane. The Details and Related Items tabs and options for the policy appear in the lower pane. On the Related Items tab, you can also apply the selected policy to object groups.

You can use the far right column of the **Active Policies** tab to reorder and therefore reprioritize the policies by dragging them to a new position. However, even though it seems like you can drag a custom policy below the default policy, you cannot. The default policy is always the last policy in the list after the view is refreshed.

How to Prioritize Policies

To set the policy priority, on the Active Policies tab, click the policy row and drag it to place it at the desired priority in the list. The priority for the Default Policy is always designated with the letter D.

Where You Manage the Active Policies

To manage the active policies, in the menu, click **Administration**, and then in the left pane click **Policies**. The **Active Policies** tab appears and lists the policies that are active for the objects in your environment.

Table 4-2. Active Policies Tab Options

Option	Description
Toolbar	<p>Use the toolbar selections to take action on the active policies.</p> <ul style="list-style-type: none"> ■ Show Association. Opens the Related Items tab so that you can associate the policy with groups. ■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to 0, which gives that policy the highest priority.
Active Policies Tab data grid	<p>vRealize Operations Manager displays the priority and high-level details for the active policies.</p> <ul style="list-style-type: none"> ■ Priority. Ranking of the priority of the policy. The default policy is marked with a check mark in the Is Default column. ■ Name. Name of the policy as it appears in the Add or Edit Monitoring Policy wizard, and in areas where the policy applies to objects, such as in Custom Groups. ■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects. ■ Groups. Indicates the number of object groups to which the policy is assigned. ■ Affected Objects. Displays the object name, type, and adapter to which the active policy is assigned, and the direct parent group, when applicable. ■ Last Modified. Date and time that the policy was last modified. ■ Modified By. User who last modified the policy settings.

Table 4-2. Active Policies Tab Options (continued)

Option	Description
Active Policies Tab > Details Tab	<p>The Details tab displays the name and description of the policy from which the settings are inherited, the policy priority, who last modified the policy, and the number of object groups associated with the policy. From the Details tab, you can view the settings that are locally defined in your policy, and the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.</p> <ul style="list-style-type: none"> ■ Locally Defined Settings. Displays the locally changed policy element settings for each object type in the policy. ■ Complete Settings Including Inherited. Displays all of the policy element settings for each object type in the policy, including locally changed settings and settings that are inherited. A summary of the enabled and disabled alert definitions, symptom definitions, and attributes appear indicate the number of changes in the policy. The policy element settings include symptom thresholds, and indicate changes made to the Workload, Capacity Remaining, and Time Remaining settings.
Active Policies Tab > Related Objects Tab	<p>Summarizes the related groups and objects, and details about the selected object group and objects.</p> <ul style="list-style-type: none"> ■ Groups. Displays the groups of objects associated with the selected active policy, and provides options to add and release an association. <ul style="list-style-type: none"> ■ Add Association. Opens the Apply the policy to groups dialog box where you select object groups to associate with the selected policy. ■ Release Association. Opens a confirmation dialog box to confirm the release of the object group that is associated with the selected policy. ■ Data grid. Displays the groups assigned to this policy, the object types associated with the group, and the number of objects in the group. ■ Details for the selected object group. Displays the object group name, type, and number of members associated with the selected policy, and the type of association with the policy. An object group can have a direct association with a policy, and inherited policy associations based on the base policies that you selected when you created a local policy. For example, if the Base Settings policy appears in the list, with an inherited association, the Base Settings policy was included in the base policies selected when this policy was created. ■ Affected Objects. Displays the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it appears in this data grid.

Policy Library Tab for Policies

The **Policy Library** tab displays the base settings, default policy, and other best practice policies that vRealize Operations Manager includes. You can use the library policies to create your own policies. The policy library includes all the configurable settings for the policy elements, such as workload, capacity and time remaining, and so on.

How the Policy Library Works

Use the options on the **Policy Library** tab to create your own policy from an existing policy, or to override the settings from an existing policy so that you can apply the new settings to groups of objects. You can also import and export a policy.

To display the details for a selected policy, click the split bar to expand the pane. The Details and Related Items tabs and options for the policy appear in the lower pane. On the Related Items tab, you can also apply the selected policy to object groups.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can also apply the policy to object groups. To update the policy association with an object group, the role assigned to your user account must have the Manage Association permission enabled for policy management.

Where You Manage the Policy Library

To manage the policy library, in the menu, click **Administration**, and then in the left pane click **Policies**. The **Policy Library** tab appears and lists the policies available to use for your environment.

Table 4-3. Policy Library Tab Options

Option	Description
Toolbar	<p>Use the toolbar selections to take action in the policy library.</p> <ul style="list-style-type: none"> ■ Add New Policy. Create a policy from an existing policy. ■ Edit Selected Policy. Customize the policy so that you can override settings for vRealize Operations Manager to analyze and report data about the associated objects. ■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to 0, which gives that policy the highest priority. ■ Import Policy and Export Policy. You can import or export a policy in XML format. To import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management. ■ Delete Selected Policy. Remove a policy from the list.
Policy Library Tab data grid	<p>vRealize Operations Manager displays the high-level details for the policies.</p> <ul style="list-style-type: none"> ■ Name. Name of the policy as it appears in the Add or Edit Monitoring Policy wizard, and in areas where the policy applies to objects, such as in Custom Groups. ■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects. ■ Last Modified. Date and time that the policy was last modified. ■ Modified By. User who last modified the policy settings.

Table 4-3. Policy Library Tab Options (continued)

Option	Description
Policy Library Tab > Details Tab	<p>The Details tab displays the name and description of the policy from which the settings are inherited, the policy priority, who last modified the policy, and the number of object groups associated with the policy. From the Details tab, you can view the settings that are locally defined in your policy, and the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.</p> <ul style="list-style-type: none"> ■ Locally Defined Settings. Displays the locally changed policy element settings for each object type in the policy. ■ Complete Settings Including Inherited. Displays all of the policy element settings for each object type in the policy, including locally changed settings and settings that are inherited. A summary of the enabled and disabled alert definitions, symptom definitions, and attributes appear indicate the number of changes in the policy. The policy element settings include symptom thresholds, and indicate changes made to the Workload, Capacity Remaining, and Time Remaining settings.
Related Objects Tab	<p>Summarizes the related groups and objects, and details about the selected object group and objects.</p> <ul style="list-style-type: none"> ■ Groups. Displays the groups of objects associated with the selected active policy, and provides options to add and release an association. <ul style="list-style-type: none"> ■ Add Association. Opens the Apply the policy to groups dialog box where you select object groups to associate with the selected policy. ■ Release Association. Opens a confirmation dialog box to confirm the release of the object group that is associated with the selected policy. ■ Data grid. Displays the groups assigned to this policy, the object types associated with the group, and the number of objects in the group. ■ Details for the selected object group. Displays the object group name, type, and number of members associated with the selected policy, and the type of association with the policy. An object group can have a direct association with a policy, and inherited policy associations based on the base policies that you selected when you created a local policy. For example, if the Base Settings policy appears in the list, with an inherited association, the Base Settings policy was included in the base policies selected when this policy was created. ■ Affected Objects. Displays the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it appears in this data grid.

Operational Policies

Determine how to have vRealize Operations Manager monitor your objects, and how to notify you about problems that occur with those objects.

vRealize Operations Manager Administrators assign policies to object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Enable and disable alerts.

- Control data collections by persisting or not persisting metrics on the objects in your environment.
- Configure the product analytics and thresholds.
- Monitor objects and applications at different service levels.
- Prioritize policies so that the most important rules override the defaults.
- Understand the rules that affect the analytics.
- Understand which policies apply to object groups.

vRealize Operations Manager includes a library of built-in active policies that are already defined for your use. vRealize Operations Manager applies these policies in priority order.

When you apply a policy to an object group, vRealize Operations Manager collects data from the objects in the object group based on the thresholds, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that are enabled in the policy.

The following examples of policies might exist for a typical IT environment.

- Maintenance: Optimized for ongoing monitoring, with no thresholds or alerts.
- Critical Production: Production environment ready, optimized for performance with sensitive alerting.
- Important Production: Production environment ready, optimized for performance with medium alerting.
- Batch Workloads: Optimized to process jobs.
- Test, Staging, and QA: Less critical settings, fewer alerts.
- Development: Less critical settings, no alerts.
- Low Priority: Ensures efficient use of resources.
- Default Policy: Default system settings.

Types of Policies

There are three types of policies such as default policies, custom policies, and policies that are offered with vRealize Operations Manager.

Custom Policies

You can customize the default policy and base policies included with vRealize Operations Manager for your own environment. You can then apply your custom policy to groups of objects, such as the objects in a cluster, or virtual machines and hosts, or to a group that you create to include unique objects and specific criteria.

You must be familiar with the policies so that you can understand the data that appears in the user interface, because policies drive the results that appear in the vRealize Operations Manager dashboards, views, and reports.

To determine how to customize operational policies and apply them to your environment, you must plan ahead. For example:

- Must you track CPU allocation? If you overallocate CPU, what percentage must you apply to your production and test objects?
- Will you overallocate memory or storage? If you use High Availability, what buffers must you use?
- How do you classify your logically defined workloads, such as production clusters, test or development clusters, and clusters used for batch workloads? Or, do you include all clusters in a single workload?
- How do you capture peak use times or spikes in system activity? In some cases, you might need to reduce alerts so that they are meaningful when you apply policies.

When you have privileges applied to your user account through the roles assigned, you can create and modify policies, and apply them to objects. For example:

- Create a policy from an existing base policy, inherit the base policy settings, then override specific settings to analyze and monitor your objects.
- Use policies to analyze and monitor vCenter Server objects and non-vCenter Server objects.
- Set custom thresholds for analysis settings on all object types to have vRealize Operations Manager report on workload, and so on.
- Enable specific attributes for collection, including metrics, properties, and super metrics.
- Enable or disable alert definitions and symptom definitions in your custom policy settings.
- Apply the custom policy to object groups.

When you use an existing policy to create a custom policy, you override the policy settings to meet your own needs. You set the allocation and demand, the overcommit ratios for CPU and memory, and the thresholds for capacity risk and buffers. To allocate and configure what your environment is actually using, you use the allocation model and the demand model together. Depending on the type of environment you monitor, such as a production environment versus a test or development environment, whether you over allocate at all and by how much depends on the workloads and environment to which the policy applies. You might be more conservative with the level of allocation in your test environment and less conservative in your production environment.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Your policies are unique to your environment. Because policies direct vRealize Operations Manager to monitor the objects in your environment, they are read-only and do not alter the state of your objects. For this reason, you can override the policy settings to fine-tune them until vRealize Operations Manager displays the results that are meaningful and that affect for your environment. For example, you can adjust the capacity buffer settings in your policy, and then view the data that appears in the dashboards to see the effect of the policy settings.

Default Policy in vRealize Operations Manager

The default policy is a set of rules that applies to the majority of your objects.

The Default policy appears on the **Active Policies** tab, and is marked with the letter D in the Priority column. The Default policy can apply to any number of objects.

The Default policy always appears at the bottom in the list of policies, even if that policy is not associated with an object group. When an object group does not have a policy applied, vRealize Operations Manager associates the Default policy with that group.

A policy can inherit the Default policy settings, and those settings can apply to various objects under several conditions.

The policy that is set to Default always takes the lowest priority. If you attempt to set two policies as the Default policy, the first policy that you set to Default is initially set to the lowest priority. When you set the second policy to Default, that policy then takes the lowest priority, and the earlier policy that you set to Default is set to the second lowest priority.

You can use the Default policy as the base policy to create your own custom policy. You modify the default policy settings to create a policy that meets your analysis and monitoring needs. When you start with the Default policy, your new policy inherits all of the settings from the Default base policy. You can then customize your new policy and override these settings.

The data adapters and solutions installed in vRealize Operations Manager provide a collective group of base settings that apply to all objects. In the policy navigation tree on the **Policy Library** tab, these settings are called Base Settings. The Default policy inherits all of the base settings by default.

Policies Provided with vRealize Operations Manager

vRealize Operations Manager includes sets of policies that you can use to monitor your environment, or as the starting point to create your own policies.

Verify that you are familiar with the policies provided with vRealize Operations Manager so that you can use them in your own environment, and to include settings in new policies that you create.

Where You Find the Policies Provided with vRealize Operations Manager Policies

In the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. To see the policies provided with vRealize Operations Manager, expand the Base Settings policy.

Policies That vRealize Operations Manager Includes

All policies exist under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects. In the policy navigation tree on the **Policy Library** tab, these settings are called Base Settings.

The Base Settings policy is the umbrella policy for all other policies, and appears at the top of the policy list in the policy library. All of the other policies reside under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects.

The Config Wizard Based Policy set includes policies provided with vRealize Operations Manager that you use for specific settings on objects to report on your objects. The Config Wizard Based Policy set includes several types of policies:

- Efficiency alerts policies for infrastructure objects and virtual machines
- Health alerts policies for infrastructure objects
- Overcommit policies for CPU and Memory
- Risk alerts policies for infrastructure objects and virtual machines

The Default Policy includes a set of rules that applies to the majority of your objects.

The VMware Management Policies set includes policies that you use for your type of environment, such as production as opposed to test and development. These policies contain settings that monitor for peak periods, batch and interactive workloads, and demand and allocation models. The VMware Management Policies set provided with vRealize Operations Manager include the following policies:

Table 4-4. Functions of VMware Management Policies

VMware Management Policy	What it does
VMware Excludes over-sized analysis	Does not calculate reclaimable capacity from oversized virtual machines
VMware Optimized for 15-minute peak periods	Configured to cause capacity alerts for workloads that spike for 15 minutes.
VMware Optimized for 30-minute peak periods	Configured to cause capacity alerts for workloads that spike for 30 minutes.
VMware Policy for Batch workloads	Optimized for batch workloads that run less than four hours.

Table 4-4. Functions of VMware Management Policies (continued)

VMware Management Policy	What it does
VMware Policy for Interactive workloads	Configured to be sensitive toward interactive workloads, such as a desktop or Web server, based on 15-minute peaks with large buffers.
VMware Production Policy (Demand only)	Optimized for production loads, without using allocation limits, to obtain the most capacity.
VMware Production Policy (with Allocation)	Optimized for production loads that require the demand and allocation capacity models.
VMware Production Policy (without Allocation)	Optimized for production loads that require demand capacity models, and provides the highest overcommit without contention.
VMware Test and Dev Policy (without Allocation).	Optimized for Dev and Test environments to maximize capacity without causing significant contention, because it does not include capacity planning at the virtual machine level.

Using the Monitoring Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the monitoring policy workspace to create local policies quickly, and update the settings in existing policies. Select a base policy to use as the source for your local policy settings, and modify the thresholds and settings used for analysis and collection of data from groups of objects in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated object groups.

Prerequisites

Verify that objects groups exist for vRealize Operations Manager to analyze and collect data, and if they do not exist, create them. See [Managing Custom Object Groups in VMware vRealize Operations Manager](#).

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Policies**.
- 2 Click **Policy Library**, and click the **Add New Policy** icon to add a policy, or select the policy and click the **Edit Selected Policy** icon to edit an existing policy.

You can add and edit policies on the **Policy Library** tab, and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.
- 3 In the Getting Started workspace, assign a name and description to the policy.

Give the policy a meaningful name and description so that all users know the purpose of the policy.

- 4 Click **Select Base Policies**, and in the workspace, select one or more policies to use as a baseline to define the settings for your new local policy.

When you create a policy, you can use any of the policies provided with vRealize Operations Manager as a baseline source for your new policy settings.

- 5 Click **Override Settings**, and in the workspace, filter the object types to customize your policy for the objects to associate with this policy.

Filter the object types, and modify the settings for those object types so that vRealize Operations Manager collects and displays the data that you expect in the dashboards and views.

- 6 Click **Override Attributes**, and in the workspace, select the metric, property, or super metric attributes to include in your policy.

vRealize Operations Manager collects data from the objects in your environment based on the metric, property, or super metric attributes that you include in the policy.

- 7 Click **Override Alert / Symptom Definitions**, and in the workspace, enable or disable the alert definitions and symptom definitions for your policy.

vRealize Operations Manager identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

- 8 Click **Apply Policy to Groups**, and in the workspace, select one or more groups to which the policy applies.

VMware vRealize Operations Manager monitors the objects according to the settings in the policy that is applied to the object group, triggers alerts when thresholds are violated, and reports the results in the dashboards, views, and reports. If you do not assign a policy to one or more object groups, VMware vRealize Operations Manager does not assign the settings in that policy to any objects, and the policy is not active. For an object group that does not have a policy assigned, VMware vRealize Operations Manager associates the object group with the Default Policy.

- 9 Click **Save** to retain the settings defined for your local policy.

What to do next

After vRealize Operations Manager analyzes and collects data from the objects in your environment, review the data in the dashboards and views. If the data is not what you expected, edit your local policy to customize and override the settings until the dashboards display the data that you need.

Policy Workspace in vRealize Operations Manager

The policy workspace allows you to quickly create and modify policies. To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

How the Policy Workspace Works

Every policy includes a set of packages, and uses the defined problems, symptoms, metrics, and properties in those packages to apply to specific object groups in your environment. You can view details for the settings inherited from the base policy, and display specific settings for certain object types. You can override the settings of other policies, and include additional policy settings to apply to object types.

Use the **Add** and **Edit** options to create policies and edit existing policies.

Where You Create and Modify a Policy

To create and modify policies, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. The policy workspace is where you select the base policies, and customize and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can apply the policy to object groups.

To remove a policy from the list, select the policy and click the red X.

Policy Workspace Options

The policy workspace includes a step-by-step workflow to create and edit a policy, and apply the policy to custom object groups.

- [Getting Started Details](#)

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

- [Select Base Policy Details](#)

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a new policy. In the policy content area, you can view the packages and elements for the base policy and additional policies that you selected to override the settings, and compare the differences in settings highlighted between these policies. You select the settings and objects types to display.

- [Analysis Settings Details](#)

You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

- [Workload Automation Details](#)

You can set the workload automation options for your policy, so that vRealize Operations Manager can optimize the workload in your environment per your definition.

■ [Collect Metrics and Properties Details](#)

You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

■ [Alert and Symptom Definitions Details](#)

You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

■ [Apply Policy to Groups Details](#)

You can assign your local policy to one or more groups of objects to have VMware vRealize Operations Manager analyze those objects according to the settings in your policy, trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

Getting Started Details

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

Where You Assign the Policy Name and Description

To add a name and description to a policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Getting Started**. The name and description appear in the workspace.

Table 4-5. Name and Description Options in the Add or Edit Monitoring Policy Workspace

Option	Description
Name	Name of the policy as it appears in the Add or Edit Monitoring Policy wizard, and in areas where the policy applies to objects, such as Custom Groups.
Description	Meaningful description of the policy. For example, use the description to indicate which policy is inherited, and any specific information that users need to understand the relationship of the policy to one or more groups of objects.
Start with	<p>The base policy that will be used as a starting point. All settings from the base policy will be inherited as default settings in your new policy. You can override these settings to customize the new policy.</p> <p>Select a base policy to inherit the base policy settings as a starting point for your new policy.</p>

Select Base Policy Details

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a new policy. In the policy content area, you can view the packages and elements for the base policy and additional policies that you selected to

override the settings, and compare the differences in settings highlighted between these policies. You select the settings and objects types to display.

How the Select Base Policies Workspace Works

To create a policy, select a base policy from which your new custom policy inherits settings. To override some of the settings in the base policy according to the requirements for the service level agreement for your environment, you can select and apply a separate policy for a management pack solution. The override policy includes specific settings defined for the types of objects to override, either manually or that an adapter provides when it is integrated with vRealize Operations Manager. The settings in the override policy overwrite the settings in the base policy that you selected.

When you select and apply a policy in the left pane to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the applied policy history list in the right pane.

The right pane displays tabs for the inherited policy configuration, and your policy, and displays a preview of the selected policy tab in the Policy Preview pane. When you select one of the policy tabs, you can view the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of enabled and disabled changes.

In the right pane, you select the objects to view so that you can see which policy elements apply to the object type. For example, when you select the StorageArray object type, and you click the tab to display the configuration settings for your policy, the Policy Preview pane displays the local packages for the policy and the object group types with the number of policy elements in each group.

You can preview the policy settings for all object types, only the object types that have settings changed locally, or settings for new object types that you add to the list, such as Storage Array storage devices.

Where You Select and Override Base Policies Settings

To select a base policy to use as a starting point for your own policy, and to select a policy to override one or more settings that your policy inherits from the base policy, in the menu, select **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left add a name for the policy and click **Select Base Policy**. The policy configuration, objects, and preview appear in the workspace.

Table 4-6. Base Policy and Override Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Show changes for	<p>Select the objects to view changes.</p> <ul style="list-style-type: none"> ■ All object types. Displays the number of enabled and disabled alert definitions, symptom definitions, and metrics and properties, the number of enabled and disabled changes, and the object type groups and the number of local policy elements for each group. ■ All object types with overrides. Displays the object types that have changes applied, with the objects types selected for override. Use the drop-down menu to select object types. Click the filter button to add the selected object type to the list so that you can preview and configure the settings. ■ Add settings for new set of objects. Provides a list of the object types so that you can select an object type, such as Storage Devices > SAN, and add the selected object to the Object types list.
Override settings from additional policies	Select and apply one or more policies to override the settings that your policy inherits from the base policy.
Apply	Applies the override policy to your policy, and lists the override policy in the applied policy history.
Applied policy template history	Displays the policies that you selected to override the settings in your policy.
Configuration inherited from base policy	When selected, displays a preview of the inherited policy configuration in the Policy Preview pane.
Configuration settings defined in this policy	When selected, displays a preview of your policy configuration in the Policy Preview pane.
Policy Preview	<p>Displays summary information about the local packages and object group types.</p> <ul style="list-style-type: none"> ■ Packages (Local). Displays the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of policy elements for each object group. ■ Object Type groups. Displays the associated object groups. ■ Drop down arrows on packages and settings. Displays the packages and settings for the displayed policies.

Analysis Settings Details

You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

How the Analysis Settings Workspace Works

When you turn on and configure the analysis settings for a policy, you can override the settings for the policy elements that vRealize Operations Manager uses to trigger alerts and display data. These types of settings include symptom thresholds based on alerts, situational settings such as committed projects to calculate capacity and time remaining, and other detailed settings.

You expand a policy element setting and configure the values to make your policy specific. For example, to reclaim capacity, you can set percentages to have vRealize Operations Manager indicate when a resource is oversized, idle, or powered off.

Policies focus on objects and object groups. When you configure policy element settings for your local policy, you must consider the object type and the results that you expect to see in the dashboards and views. If you do not make any changes to the settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

Where You Set the Policy Analysis Settings

To set the analysis settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Analysis Settings**. The analysis settings for host systems, virtual machines, and other object types that you select appear in the workspace.

Table 4-7. Analysis Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Show changes for	<p>Select the objects to view changes.</p> <ul style="list-style-type: none"> ■ All object types. Displays the number of enabled and disabled alert definitions, symptom definitions, and metrics and properties, the number of enabled and disabled changes, and the object type groups and the number of local policy elements for each group. ■ All object types with overrides. Displays the object types that have changes applied, with the objects types selected for override. Use the drop-down menu to select object types. Click the filter button to add the selected object type to the list so that you can preview and configure the settings. ■ Add settings for new set of objects. Provides a list of the object types so that you can select an object type, such as Storage Devices > SAN, and add the selected object to the Object types list.
Right pane - Analysis Settings for object types	<p>The right pane displays a list of the object types that you selected in the left pane. Expand a view of the policy elements and settings for the object type so that you can have vRealize Operations Manager analyze the object type.</p> <p>Expand the view for the object type so that you can view and modify the threshold settings for the following policy elements:</p> <ul style="list-style-type: none"> ■ Workload ■ Time Remaining ■ Capacity Remaining ■ Compliance ■ Maintenance Schedule <p>Click the lock icon on the right of each element to override the settings and change the thresholds for your policy.</p>
Time Remaining Calculations	<p>You can set the risk level for the time that is remaining when the forecasted total need of a metric reaches usable capacity.</p> <ul style="list-style-type: none"> ■ Conservative. Select this option for production and mission critical workloads. ■ Aggressive. Select this option for non-critical workloads.

Policy Workload Element

Workload is a measurement of the demand for resources on an object. You can turn on and configure the settings for the Workload element for the object types in your policy.

How the Workload Element Works

The Workload element determines how vRealize Operations Manager reports on the resources that the selected object group uses. The resources available to the object group depend on the amount of configured and usable resources.

- A specific amount of physical memory is a configured resource for a host system, and a specific number of CPUs is a configured resource for a virtual machine.
- The usable resource for an object or an object group is a subset of, or equal to, the configured amount.
- The configured and usable amount of a resource can vary depending on the type of resource and the amount of virtualization overhead required, such as the memory that an ESX host machine requires to run the host system. When accounting for overhead, the resources required for overhead are not considered to be usable, because of the reservations required for virtual machines or for the high availability buffer.

Where You Override the Policy Workload Element

To view and override the policy workload analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The workload settings for the object types that you selected appear in the right pane.

View the Workload policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-8. Policy Workload Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Workload Score Threshold	Allows you to set the number of collection cycles it takes to trigger or clear an alert.

Policy Time Remaining Element

The Time remaining element is a measure of the amount of time left before your objects run out of capacity.

How the Time Remaining Element Works

The Time Remaining element determines how vRealize Operations Manager reports on the available time until capacity runs out for a specific object type group.

- The time remaining indicates the amount of time that remains before the object group consumes the capacity available. vRealize Operations Manager calculates the time remaining as the number of days remaining until all the capacity is consumed.

- To keep the Time Remaining more than the critical threshold setting or to keep it green, your objects must have more days of capacity available.

Where You Override the Policy Time Remaining Element

To view and override the policy Time Remaining analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The time remaining settings for the object types that you selected in the workspace appear in the right pane.

View the Time Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-9. Policy Time Remaining Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Time Remaining Score Threshold	Allows you to set the number of days until capacity is projected to run out based on your current consumption trend.

Policy Capacity Remaining Element

Capacity is a measurement of the amount of memory, CPU, and disk space for an object. You can turn on and configure the settings for the Capacity Remaining element for the object types in your policy.

How the Capacity Remaining Element Works

The Capacity Remaining element determines how vRealize Operations Manager reports on the available capacity until resources run out for a specific object type group.

- The capacity remaining indicates the capability of your environment to accommodate workload.
- Usable capacity is a measurement of the percentage of capacity available, minus the capacity affected when you use high availability.

Where You Override the Policy Capacity Remaining Element

To view and override the policy Capacity Remaining analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The capacity remaining settings for the object types that you selected in the workspace appear in the right pane.

View the Capacity Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-10. Policy Capacity Remaining Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Capacity Remaining Score Threshold	Allows you to set the percentage at which the capacity remaining alerts must be triggered.

Policy Compliance Element

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the Compliance element for the object types in your policy.

Where You Override the Policy Compliance Element

To view and override the policy Compliance analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The compliance settings for the object types that you selected appear in the right pane.

View the Compliance policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-11. Policy Compliance Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Compliance Score Threshold	Allows you to set the compliance score threshold based on the number of violations against those standards.

Policy Maintenance Schedule Element

You can set a time to perform maintenance tasks for each policy.

Where You Override the Policy Maintenance Schedule Element

To view and override the policy Maintenance Schedule analysis setting, in the menu, click **Administration**, and then in the left pane, click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The maintenance schedule settings for the object types that you selected appear in the right pane.

View the maintenance schedule policy element.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-12. Policy Maintenance Schedule Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Maintenance Schedule	Sets a time to perform maintenance tasks. During maintenance, vRealize Operations Manager does not calculate analytics.

Workload Automation Details

You can set the workload automation options for your policy, so that vRealize Operations Manager can optimize the workload in your environment per your definition.

How the Workload Automation Workspace Works

You click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

Where You Set the Policy Workload Automation

To set the workload automation for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Workload Automation**.

Table 4-13. Workload Automation in the Add or Edit Monitoring Policy Workspace

Option	Description
Workload Optimization	<p>Select a goal for workload optimization.</p> <p>Select Balance when workload performance is your first goal. This approach proactively moves workloads so that the resource utilization is balanced, leading to maximum headroom for all resources.</p> <p>Select Moderate when you want to minimize the workload contention.</p> <p>Select Consolidate to proactively minimize the number of clusters used by workloads. You might be able to repurpose resources that are freed up. This approach is good for cost optimization, while making sure that performance goals are met. This approach might reduce licensing and power costs.</p>
Cluster Headroom	Headroom establishes a required capacity buffer, for example, 20 percent. It provides you with an extra level of control and ensures that you have extra space for growth inside the cluster when required. Defining a large headroom setting limits the systems opportunities for optimization.
Advanced Settings	Click Advanced Settings to select what type of virtual machines vRealize Operations Manager moves first to address workload. You can set Storage vMotion on or off. The default is ON.

Collect Metrics and Properties Details

You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

How the Collect Metrics and Properties Workspace Works

When you create or customize a policy, you can override the base policy settings to have vRealize Operations Manager collect the data that you intend to use to generate alerts, and report the results in the dashboards.

To define the metric and super metric symptoms, metric event symptoms, and property symptoms, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Symptom Definitions**.

Where You Override the Policy Attributes

To override the attributes and properties settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Collect Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

Table 4-14. Collect Metrics and Properties Options






Option	Description
Actions	Select one or more attributes and select enable, disable, or inherit to change the state and KPI for this policy.
Filter options	<p>Deselect the options in the Attribute Type, State, KPI, and DT drop-down menus, to narrow the list of attributes.</p> <ul style="list-style-type: none"> ■  Enabled. Indicates that an attribute will be calculated. ■  Enabled (Force). Indicates state change due to a dependency. ■  Disabled. Indicates that an attribute will not be calculated. ■  Inherited. Indicates that the state of this attribute is inherited from the base policy and will be calculated. ■  Inherited. Indicates that the state of this attribute is inherited from the base policy and will not be calculated. <p>The KPI determines whether the metric, property, or super metric attribute is considered to be a key performance indicator (KPI) when vRealize Operations Manager reports the collected data in the dashboards. Filter the KPI states to display attributes with KPI enabled, disabled, or inherited for the policy.</p>
Object Type	Filters the attributes list by object type.

Table 4-14. Collect Metrics and Properties Options (continued)

Option	Description
Page Size	The number of attributes to list per page.
Attributes data grid	<p>Display the attributes for a specific object type.</p> <ul style="list-style-type: none"> ■ Name. Identifies the name of the metric or property for the selected object type. ■ Type. Distinguishes the type of attribute to be either a metric, property, or super metric. ■ Adapter Type. Identifies the adapter used based on the object type selected, such as Storage Devices. ■ Object Type. Identifies the type of object in your environment, such as StorageArray. ■ State. Indicates whether the metric, property, or super metric is inherited from the base policy. ■ KPI. Indicates whether the key performance indicator is inherited from the base policy. If a violation against a KPI occurs, vRealize Operations Manager generates an alert. ■ DT. Indicates whether the dynamic threshold (DT) is inherited from the base policy.

Alert and Symptom Definitions Details

You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

How the Alert and Symptom Definitions Workspace Works

vRealize Operations Manager collects data for objects and compares the collected data to the alert definitions and symptom definitions defined for that object type. Alert definitions include associated symptom definitions, which identify conditions on attributes, properties, metrics, and events.

You can configure your local policy to inherit alert definitions from the base policies that you select, or you can override the alert definitions and symptom definitions for your local policy.

Before you add or override the alert definitions and symptom definitions for a policy, familiarize yourself on the available alerts and symptoms.

- To view the available alert definitions, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Alert Definitions**.
- To view the available symptom definitions, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Symptom Definitions**. Symptom definitions are available for metrics, properties, messages, faults, smart early warnings, and external events.

A summary of the number of problem and symptoms that are enabled and disabled, and the difference in changes of the problem and symptoms as compared to the base policy, appear in the Analysis Settings pane of the policies workspace.

Where You Override the Alert Definitions and Symptom Definitions

To override the alert definitions and symptom definitions for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Alert / Symptom Definitions**. The definitions appear in the workspace.

Policy Alert Definitions and Symptom Definitions

You can override the alert definitions and symptom definitions for each policy.

- **Policy Alert Definitions**

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

- **Policy Symptom Definitions**

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

How the Policy Alert Definitions Work

vRealize Operations Manager uses problems to trigger alerts. A problem manifests when a set of symptoms exists for an object, and requires you to take action on the problem. Alerts indicate problems in your environment. vRealize Operations Manager generates alerts when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert occurs, vRealize Operations Manager presents the triggering symptoms for you to take action.

Some of the alert definitions include predefined symptoms. When you include symptoms in an alert definition, and enable the alert, an alert is generated when the symptoms are true.

The Alert Definitions pane displays the name of the alert, the number of symptoms defined, the adapter, object types such as host or cluster, and whether the alert is enabled as indicated by **Local**, disabled as indicated by **not Local**, or inherited. Alerts are inherited with a green checkmark by default, which means that they are enabled.

You can automate an alert definition in a policy when the highest priority recommendation for the alert has an associated action.

To view a specific set of alerts, you can select the badge type, criticality type, and the state of the alert to filter the view. For example, you can set the policy to send fault alerts for virtual machines.

Where You Modify the Policy Alert Definitions

To modify the alerts associated with policies, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Alert / Symptom Definitions**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-15. Alert Definitions in the Add or Edit Monitoring Policy Workspace

Option	Description
Actions	Select one or more alert definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	<p>Deselect the options in the Type and State drop-down menus, to narrow the list of symptom definitions.</p> <p>Impact indicates the health, risk, and efficiency badges to which the alerts apply.</p> <p>Criticality indicates the information, critical, immediate, warning, or automatic criticality types to which the alert definition applies.</p> <p>Automate indicates the actions that are enabled for automation when an alert triggers, or actions that are disabled or inherited. Actions that are enabled for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to Local with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark.</p>
Object Type	Filters the alert definitions list by object type.
Page Size	The number of alert definitions to list per page.
Filter	Locates data in the alert definition list.
Alert Definitions data grid	<p>Displays information about the alert definitions for the object types. The full name for Alert definition and the criticality icon appear in a tooltip when you hover the mouse over the Alert Definition name.</p> <ul style="list-style-type: none"> ■ Name. Meaningful name for the alert definition. ■ Symptom Definitions. Number of symptoms defined for the alert. ■ Actionable Recommendations. Only recommendations with actions in the first priority, as they are the only ones you can automate. ■ Automate. When the action is set to Local, the action is enabled for automation when an alert triggers. Actions that are enabled for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to Local with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ State. Alert definition state, either enabled as indicated by Local, disabled as indicated by not Local, or inherited from the base policy.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

How the Policy Symptom Definitions Work

vRealize Operations Manager uses symptoms that are enabled to generate alerts. When the symptoms used in an alert definition are true, and the alert is enabled, an alert is generated.

When a symptom exists for an object, the problem exists and requires that you take action to solve it. When an alert occurs, vRealize Operations Manager presents the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

To assess objects for symptoms, you can include symptoms packages in your policy for metrics and super metrics, properties, message events, and faults. You can enable or disable the symptoms to determine the criteria that the policy uses to assess and evaluate the data collected from the objects to which the policy applies. You can also override the threshold, criticality, wait cycles, and cancel cycles.






The Symptoms pane displays the name of the symptom, the associated management pack adapter, object type, metric or property type, a definition of the trigger such as for CPU usage, the state of the symptom, and the trigger condition. To view a specific set of symptoms in the package, you can select the adapter type, object type, metric or property type, and the state of the symptom.

When a symptom is required by an alert, the state of the symptom is enabled, but is dimmed so that you cannot modify it. The state of a required symptom includes an information icon that you can hover over to identify the alert that required this symptom.

Where You Modify the Policy Symptom Definitions

To modify the policy package of symptoms, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left, click **Alert / Symptom Definitions**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-16. Symptom Definitions in the Add or Edit Monitoring Policy Workspace

Option	Description
Actions	Select one or more symptom definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	<p>Deselect the options in the Type and State drop-down menus, to narrow the list of symptom definitions.</p> <ul style="list-style-type: none"> ■  Enabled. Indicates that a symptom definition will be included. ■  Enabled (Force). Indicates state change due to a dependency. ■  Disabled. Indicates that a symptom definition not be included. ■  Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will be included. ■  Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will not be included. <p>Type determines whether symptom definitions that apply to HT and DT metrics, properties, events such as message, fault, and metric, and smart early warnings appear in the list.</p> <p>State determines whether enabled, disabled, and inherited symptom definitions appear in the symptom definition list.</p>
Object Type	Filters the symptom definitions list by object type
Page Size	The number of symptom definitions to list per page.
Filter	Locate data in the symptom definition list.
Symptom Definitions data grid	<p>Displays information about the symptom definitions for the object types. The full name for Symptom Definition appears in a tooltip when you hover the mouse over the Symptom Definition name.</p> <ul style="list-style-type: none"> ■ Name. Symptom definition name as defined in the list of symptom definitions in the Content area. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ Type. Object type on which the symptom definition must be evaluated. ■ Trigger. Static or dynamic threshold, based on the number of symptom definitions, the object type and metrics selected, the numeric value assigned to the symptom definition, the criticality of the symptom, and the number of wait and cancel cycles applied to the symptom definition. ■ State. Symptom definition state, either enabled, disabled, or inherited from the base policy. ■ Condition. Enables action on the threshold. When set to Override, you can change the threshold. Otherwise set to default. ■ Threshold. To change the threshold, you must set the State to Enabled, set the condition to Override, and set the new threshold in the Override Symptom Definition Threshold dialog box.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Apply Policy to Groups Details

You can assign your local policy to one or more groups of objects to have VMware vRealize Operations Manager analyze those objects according to the settings in your policy, trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

How the Apply Policy to Groups Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more object groups. VMware vRealize Operations Manager uses the settings in the policy to analyze and collect data from the associated objects, and displays the data in dashboards, views, and reports.

Where You Apply a Policy to Groups

To apply the policy to object groups, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Apply Policy to Groups**.

Apply Policy to Groups Options

To apply the policy to groups of objects, select the check box for the object group in the workspace.

You can then view the details about each object group associated with the policy. In the menu, click **Administration**, and then in the left pane click **Policies**. Click **Active Policies > Related Objects**. Click an object group in the list of groups, and view the summary in the Details pane.

For more information about how to create an object group, see [#unique_244](#).

For more information about how to create a policy, see [Policy Workspace in vRealize Operations Manager](#).

Configuring Super Metrics

5

The super metric is a mathematical formula that contains one or more metrics or properties. It is a custom metric that you design to help track combinations of metrics or properties, either from a single object or from multiple objects. If a single metric does not inform you about the behavior of your environment, you can define a super metric.

After you define it, you assign the super metric to one or more object types. This action calculates the super metric for the objects in that object type and simplifies the metrics display. For example, you define a super metric that calculates the average CPU usage on all virtual machines, and you assign it to a cluster. The average CPU usage on all virtual machines in that cluster is reported as a super metric for the cluster.

When the super metric attribute is enabled in a policy, you can also collect super metrics from a group of objects associated with a policy.

Because super metric formulas can be complex, plan your super metric before you build it. The key to creating a super metric that alerts you to the expected behavior of your objects is knowing your own enterprise and data. Use this checklist to help identify the most important aspects of your environment before you begin to configure a super metric.

Table 5-1. Designing a Super Metric Checklist

<input type="checkbox"/> Determine the objects that are involved in the behavior to track.	When you define the metrics to use, you can select either specific objects or object types. For example, you can select the specific objects VM001 and VM002, or you can select the object type virtual machine.
<input type="checkbox"/> Determine the metrics to include in the super metric.	If you are tracking the transfer of packets along a network, use metrics that refer to packets in and packets out. In another common use of super metrics, the metrics might be the average CPU usage or average memory usage of the object type you select.
<input type="checkbox"/> Decide how to combine or compare the metrics.	For example, to find the ratio of packets in to packets out, you must divide the two metrics. If you are tracking CPU usage for an object type, you might want to determine the average use. You might also want to determine what the highest or lowest use is for any object of that type. In more complex scenarios, you might need a formula that uses constants or trigonometric functions.

Table 5-1. Designing a Super Metric Checklist (continued)

<input type="checkbox"/> Decide where to assign the super metric.	You define the objects to track in the super metric, then assign the super metric to the object type that contains the objects being tracked. To monitor all the objects in a group, enable the super metric in the policy, and apply the policy to the object group.
<input type="checkbox"/> Determine the policy to which you add the super metric.	After you create the super metric, you add it to a policy. For more information, refer to Policy Workspace in vRealize Operations Manager .

What Else Can You Do with Super Metrics

- To see the super metrics in your environment, generate a system audit report. For more information, refer to the System Audit section in the Information Center.
- To create alert definitions to notify you of the performance of objects in your environment, define symptoms based on super metrics. For more information, refer to [About Metrics and Super Metrics Symptoms](#).
- Learn about the use of super metrics in policies. For more information, refer to [Policy Workspace in vRealize Operations Manager](#).
- Use OPS CLI commands to import, export, configure, and delete super metrics. For more information, refer to the OPS CLI documentation.
- To display metric-related widgets, create a custom set of metrics. You can configure one or more files that define different sets of metrics for a particular adapter and object types. This ensures that the supported widgets are populated based on the configured metrics and selected object type. For more information, refer to [Manage Metric Configuration](#).

This chapter includes the following topics:

- [Create a Super Metric](#)
- [Enhancing Your Super Metrics](#)
- [Exporting and Importing a Super Metric](#)

Create a Super Metric

Create a super metric when you want to check the health of your environment, but cannot find a suitable metric to perform the analysis.

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Configuration > Super Metrics**.
- 2 Click the **Add** icon.
- 3 Enter a meaningful name for the super metric such as **SM-AvgVMCPUUsage%** in the **Name** text box.

4 Define the formula for the super metric.

Select the function or operator. Select the metrics or attribute types that you want to use in each function or with each operator. For example, to add a super metric that captures average CPU usage across all virtual machines, perform the following tasks.

- a For Function, select **avg**.
- b In the **Operators** text box, select the left parenthesis, then select the right parenthesis. Click between the two parentheses to position the cursor.
- c Optionally, click the **This object** icon.

If the **This object** icon is selected, the super metric is calculated using the metrics and properties of the object you use for the super metric. If the icon is not selected, the super metric is calculated using metrics and properties of other objects and the super metric function displays the object with a long description.

- d In the **Adapter Type** text box of the Object Types pane, select **vCenter Adapter**.
- e From the list of object types, select **Virtual Machine**.
- f In the **Attribute Types** pane, expand All Metrics, and then expand the CPU category, scroll down, and double-click the **Usage (%)** metric.

The formula appears as a mathematical function. To view the formula in a textual format, click the **Show Formula Description** icon. If the formula syntax is wrong, an error message appears. The formula ends with `depth=1`. With `depth=1`, you assign the super metric to an object type that is one level above virtual machines in the relationship chain. The super metric appears as a metric for that object type. With `depth=2`, you assign the super metric to an object type that is two levels above virtual machines, for example a Cluster.

5 To assign the super metric to an object type at `depth=1`, type 2 instead of 1, so that `depth=2` is displayed.

6 Verify that the super metric formula has been created correctly.

- a Click the **Visualize Super Metric** icon.
- b In the Objects pane, double-click one of the objects listed.

A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

7 Click **Save**.

8 Associate the super metric with an object. vRealize Operations Manager calculates the super metric for the target objects and displays it as a metric for the object type.

- a In the Super Metrics workspace, select the super metric.
- b In the **Object Types** tab, click the **Add** icon.

- c In the Select Object Type text box, select the required object. For example, if you created your super metric for Host Systems under the vCenter Adapter, expand **vCenter Adapter**, and select **Host Systems**.
- d Click **Select**.

After one collection cycle, the super metric appears on each of the objects of the specified object type. For example, you defined the super metric to calculate average CPU usage across all virtual machines and assigned it to the Host System object type. After one collection cycle, the super metric appears as a super metric on each host.

What to do next

In the **Policies > Edit Policy > Attributes** workspace, you must select and enable each super metric. See [Custom Policies](#). Wait at least one collection cycle for the super metric to begin collecting and processing data. Then review your super metric in the **All Metrics** tab.

Enhancing Your Super Metrics

You can enhance your super metrics by using clauses and resource entry aliasing.

Where Clause

The **where** clause verifies whether a particular metric value can be used in the super metric. Use this clause to point to a different metric of the same object, such as

where = "metric_group|my_metric > 0.

For example:

```
count(${objecttype = ExampleAdapter, adaptertype = ExampleObject, metric = ExampleGroup|Rating, depth=2, where = "==1"})
```

Resource Entry Aliasing

Resource entries are used to retrieve metric data from vRealize Operations Manager for computing super metrics. A resource entry is the part of an expression which begins with **\$** followed by a **{..}** block. When computing a super metric, you might have to use the same resource entry multiple times. If you have to change your computation, you must change every resource entry, which might lead to errors. You can use resource entry aliasing to rewrite the expression.

The following example, shows a resource entry that has been used twice.

```
(min(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|active_longterm_load, depth=5, where=">=0"}) + 0.0001)/(max(${adaptertype=VMWARE, objecttype=HostSystem, attribute=cpu|demand|active_longterm_load, depth=5, where=">=0"}) + 0.0001)"
```

The following example shows how to write the expressing using resource entry aliasing. The output of both expressions is the same.


```
(min({adaptype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=">=0"} as cpuload) + 0.0001)/(max(cpuload) +
0.0001)"
```

Follow these guidelines when you use resource entry aliasing:

- When you create an alias, make sure that after the resource entry you write **as** and then **alias:name**. For example: **\${...} as alias_name**.
- The alias cannot contain the `()[]+-%/|&!=<>.,?:$` special characters, and cannot begin with a digit.
- An alias name, like all names in super metric expressions, is case-insensitive.
- Use of an alias name is optional. You can define the alias, and not use it in an expression.
- Each alias name can be used only once. For example:
\${resource1,...} as r1 + \${resource2,...} as R1.
- You can specify multiple aliases for the same resource entry. For example: **\${...} as a1 as a2.**

Conditional Expression ?: Ternary Operators

You can use a ternary operator in an expression to run conditional expressions.

For example: **expression_condition ? expression_if_true : expression_if_false**.

The result of the conditional expression is converted to a number. If the value is not 0, then the condition is assumed as true.

For example: **-0.7 ? 10 : 20** equals 10. **2 + 2 / 2 - 3 ? 4 + 5 / 6 : 7 + 8** equals 15 (7 + 8).

Depending on the condition, either **expression_if_true** or **expression_if_false** is run, but not both of them. In this way, you can write expressions such as,

\${this, metric=cpu|demandmhz} as a != 0 ? 1/a : -1. A ternary operator can contain other operators in all its expressions, including other ternary operators.

For example: **!1 ? 2 ? 3 : 4 : 5** equals 5.

Exporting and Importing a Super Metric

You can export a super metric from one vRealize Operations Manager instance and import it to another vRealize Operations Manager instance. For example, after developing a super metric in a test environment, you can export it from the test environment and import it use in a production environment.

If the super metric to import contains a reference to an object that does not exist in the target instance, the import fails. vRealize Operations Manager returns a brief error message and writes detailed information to the log file.

Procedure

1 Export a super metric.

- a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.
- b Select the super metric to export, click the **Actions** icon and select **Export Selected Super Metric** icon.

vRealize Operations Manager creates a super metric file, for example, `SuperMetric.json`.

- c Download the super metric file to your computer.

2 Import a super metric.

- a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.
- b Click the **Actions** icon and select **Import Super Metric**.
- c (Optional). If the target instance has a super metric with the same name as the super metric you are importing, you can either overwrite the existing super metric or skip the import, which is the default.

Configuring Objects

6

Using the power of object management - including metrics and alerts - you can monitor objects, applications, and systems that must stay up and running. Some metrics and alerts are prepackaged into dashboards and policies; others you combine into custom tools

vRealize Operations Manager discovers objects in your environment and makes them available to you. With the information that vRealize Operations Manager provides, you can quickly access and configure any object. For example, you can determine if a datastore is connected or providing data, or you can power on a virtual machine.

This chapter includes the following topics:

- [Object Discovery](#)

Object Discovery

Its ability to monitor and collect data on objects in your systems environment makes vRealize Operations Manager a critical tool in maintaining system uptime and ensuring ongoing good health for all system resources from virtual machines to applications to storage - across physical, virtual and cloud infrastructures.

Following are examples of objects that can be monitored.

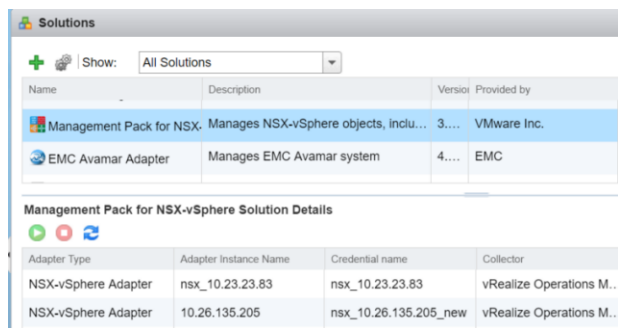
- vCenter Server
- Virtual machines
- Servers/hosts
- Compute resources
- Resource pools
- Data centers
- Storage components
- Switches
- Port groups
- Datastores

Adapters – Key to Object Discovery

vRealize Operations Manager collects data and metrics from objects using adapters, the central components of management packs, which in turn make up vRealize Operations Manager solutions. When you configure the vSphere Solution, for example, you create adapter instances customized for your environment with unique names, port numbers, and so on. You must create an adapter instance for each vCenter Server in your deployment.

Locate existing adapters in the UI as follows: in the menu, click **Administration**, then click **Solutions** in the left pane.

As shown in the screenshot, the Solutions screen lists available solutions at the top of the screen. When you select a solution, the available adapters appear in the lower half of the screen. Existing adapter instances related to each adapter are listed in the second column.



For complete information on configuring management packs and adapters, see [Chapter 1 Connecting vRealize Operations Manager to Data Sources](#)

When you create a new adapter instance, it begins discovering and collecting data from the objects designated by the adapter, and notes the relationships between them. Now you can begin to manage your objects.

About Objects

Objects are the structural components of your mission-critical IT applications: virtual machines, datastores, virtual switches and port groups are examples of objects.

Because downtime equals cost - in unused resources and lost business opportunities - it's crucial that you successfully identify, monitor and track objects in your environment. The goal is to proactively isolate, troubleshoot and correct problems even before users are aware that anything is wrong.

When a user actually reports an issue, the solution should be quick and comprehensive.

For a complete list of objects that can be defined in vRealize Operations Manager refer to [Object Discovery](#).

vRealize Operations Manager gives you visibility into objects including applications, storage and networks across physical, virtual and cloud infrastructures through a single interface that relates performance information to positive or negative events in the environment.

Managing Objects

When you monitor a large infrastructure, the number of objects and corresponding metrics in vRealize Operations Manager grows rapidly, especially as you add solutions that extend dynamic monitoring and alerts to more parts of your infrastructure. vRealize Operations Manager gives you ample tools to stay abreast of events and issues.

Adding Objects and Configuring Object Relationships

vRealize Operations Manager automatically discovers objects and their relationships once you create an adapter instance. You have the added ability to manually add any objects that you want monitored and to configure object relationships using abstract concepts rather than the connections recorded by vRealize Operations Manager. Where vRealize Operations Manager might discover the classic parent-child relationships between objects, you can create relationships between objects that might not normally be related. For example, you could configure all the datastores supporting a company department to be related.

When objects are related, a problem with one object appears as an anomaly on related objects. So object relationships can help you to identify problems in your environment quickly. The object relationships that you create are called custom groups.

Custom Groups

To create an automated management system you need some way to organize objects so that you can quickly gain insights. You can achieve a high level of automation using custom groups. You have multiple options for tailoring group attributes to support your monitoring strategy.

For example, you can designate a group either to be static or to be updated automatically with membership criteria that you designate. Consider a non-static group of all virtual machines that are powered on and have OS type Linux. When you power on a new Linux VM, it is automatically added to the group and the policy is applied.

For additional flexibility, you can also specify individual objects to be always included or excluded from a given custom group. Or you can have a different set of alerts and capacity calculations for your production environment versus your testing environments.

Managing Applications

vRealize Operations Manager allows you to create containers or objects that can contain a group of virtual machines or other objects in different structural tiers. This new application can then be managed as a single object, and have health badges and alarms aggregated from the child objects of the group.

For example, the system administrator of an online training system might request that you monitor components in the Web, application and database tiers of the training environment. You build an application that groups related training objects together in each tier. If a problem occurs with one of the objects, it is highlighted in the application display and you can investigate the source of the problem.

The Power of Object Management

Using the power of object management, including metrics and alerts - some prepackaged into dashboards and policies, others that you combine into custom monitoring tools - you'll keep a close watch on the objects, applications and systems that must stay up and running.

Managing Objects in Your Environment

An object is the individual managed item in your environment for which vRealize Operations Manager collects data, such as a router, switch, database, virtual machine, host, and vCenter Server instances.

The system requires specific information about each object. When you configure an adapter instance, vRealize Operations Manager performs object discovery to start collecting data from the objects with which the adapter communicates.

An object can be a single entity, such as a database, or a container that holds other objects. For example, if you have multiple Web servers, you can define a single object for each Web server and define a separate container object to hold all of the Web server objects. Groups and applications are types of containers.

Categorize your objects using tags, so that you can easily find, group, or filter them later. A tag type can have multiple tag values. You or vRealize Operations Manager assigns objects to tag values. When you select a tag value, vRealize Operations Manager displays the objects associated with that tag. For example, if a tag type is Lifecycle and tag values are Development, Test, Pre-production, and Production, you might assign virtual machine objects VM1, VM2, or VM3 in your environment to one or more of these tag values, depending on the virtual machine function.

Adding an Object to Your Environment

You might want to add an object by providing its information to vRealize Operations Manager. For example, some solutions cannot discover all the objects that might be monitored. For these solutions, you must either use manual discovery or manually add the object.

When you add an individual object, you provide specific information about it, including the kind of adapter to use to make the connection and the connection method. For example, a vSAN adapter does not know the location of the vSAN devices that you want to monitor.

Prerequisites

Verify that an adapter is present for the object you plan to add. See the *vRealize Operations Manager vApp Deployment and Configuration Guide*.

Procedure

- 1 In the menu, click **Administration**, then select **Configuration > Inventory Explorer** from the left pane.
- 2 On the toolbar, click the plus sign.

- 3 Use the topic menus to reveal all fields and provide the required information.

Option	Description
Display name	Enter a name for the object. For example, enter vSAN-Host1 .
Description	Enter any description. For example, enter vSAN-Host monitored with vSAN adapter
Adapter type	Select an adapter type. For example, select vSAN Adapter .
Adapter instance	Select an adapter instance.
Object type	Select an object type. For a vSAN adapter, you might select vSAN-Host. When you select the object type, the dialog box selections change to include information you provide so that vRealize Operations Manager can find and connect with the selected object type.
Host IP address	Enter the host IP. For example, enter the IP address of vSAN-Host1.
Port number	Accept the default port number or enter a new value.
Credential	Select the Credential, or click the plus sign to add new login credentials for the object.
Collection interval	Enter the collection interval, in minutes. For example, if you expect the host to generate performance data every 5 minutes, set the collection interval to 5 minutes.
Dynamic Thresholding.	Accept the default, Yes.

- 4 Click **OK** to add the object.

Results

vSAN-Host1 appears in the Inventory Explorer as a host object type for the vSAN adapter type.

What to do next

For each new object, vRealize Operations Manager assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags.

Configuring Object Relationships

vRealize Operations Manager shows the relationship between objects in your environment. Most relationships are automatically formed when the objects are discovered by an installed adapter. In addition, you can use vRealize Operations Manager to create relationships between objects that might not normally be related.

Objects are related physically, logically, or structurally.

- Physical relationships represent how objects connect in the physical world. For example, virtual machines running on a host are physically related.
- Logical relationships represent business silos. For example, all the storage objects in an environment are related to one another.
- Structural relationships represent a business value. For example, all the virtual machines that support a database are structurally related.

Solutions use adapters to monitor the objects in your environment so that physical relationship changes are reflected in vRealize Operations Manager. To maintain logical or structural relationships, you can use vRealize Operations Manager to define the object relationships. When objects are related, a problem with one object appears as an influence on related objects. So object relationships can help you to identify problems in your environment quickly.

Adding an Object Relationship

Parent-child relationships normally occur between interrelated objects in your environment. For example, a data center object for a vCenter Adapter instance might have datastore, cluster, and host system child objects.

The most common object relationships gather similar objects into groups. When you define a custom group with parent objects, a summary of that group shows alerts for that object and for any of its descendants. You can create relationships between objects that might not normally be related. For example, you might define a child object for an object in the group. You define these types of relationships by configuring object relationships.

Procedure

- 1 At the Home page, select **Administration**. Then select **Configuration > Object Relationships** in the left pane.
- 2 In the Parent Selection column, expand the object tag and select a tag value that contains the object to act as the parent object.

The objects for the tag value appear in the top pane of the second column.

- 3 Select a parent object.

Current child objects appear in the bottom pane of the second column.

- 4 In the column to the right of the List column, expand the object tag and select a tag value that contains the child object to relate to the parent.
- 5 (Optional) If the list of objects is long, filter the list to find the child object or objects.

Option	Action
Navigate the object tag list for an object	Expand the object tag in the pane to the right of the List column and select a tag value that contains the object. The objects for the tag value appear in the List column. If you select more than one value for the same tag, the list contains objects that have either value. If you select values for two or more different tags, the list includes only objects that have all of the selected values.
Search for an object by name	If you know all or part of the object name, enter it in the Search text box and press Enter.

- 6 To make an object a child object of the parent object, select the object from the list and drag it to the parent object in the top pane of the second column, or click the **Add All Objects To Parent** icon to make all of the listed objects children of the parent object.

You can use Ctrl+click to select multiple objects or Shift+click to select a range of objects.

Example: Custom Group with Child Objects

If you want vRealize Operations Manager to monitor objects in your environment to ensure that service level capacity requirements for your IT department are met, you add the objects to a custom group, apply a group policy, and define criteria that affect the membership of objects in the group. If you want to monitor the capacity of an object that does not affect the service level requirements, you can add the object as a child of a parent object in the group. If a capacity problem exists for the child object, the summary of the group shows an alert for the parent object.

Creating and Assigning Tags

A large enterprise can have thousands of objects defined in vRealize Operations Manager. Creating object tags and tag values makes it easier to find objects and metrics. With object tags, you select the tag value assigned to an object and view the list of objects that are associated with that tag value.

A tag is a type of information, for example, Adapter Types. Adapter Types is a predefined tag. Tag values are individual instances of that type of information. For example, when the system discovers objects using the vCenter Adapter, it assigns all the objects to the vCenter Adapter tag value under the Adapter Types tag.

You can assign any number of objects to each tag value, and you can assign a single object to tag values under any number of tags. You typically look for an object by looking under its adapter type, its object type, and possibly other tags.

If an object tag is locked, you cannot add objects to it. vRealize Operations Manager maintains locked object tags.

- **Predefined Object Tags**

vRealize Operations Manager includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

- **Add an Object Tag and Assign Objects to the Tag**

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

- **Use a Tag to Find an Object**

The quickest way to find an object in vRealize Operations Manager is to use tags. Using tags is more efficient than searching through the entire object list.

Predefined Object Tags

vRealize Operations Manager includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

For example, when you add an object, the system assigns it to the tag value for the collector it uses and the kind of object that it is. vRealize Operations Manager creates tag values if they do not already exist.

If a predefined tag has no values, there is no object of that tag type. For example, if no applications are defined, the applications tag has no tag values.

Each tag value appears with the number of objects that have that tag. Tag values that have no objects appear with the value zero. You cannot delete the predefined tags or tag values.

Table 6-1. Predefined Tags

Tag	Description
Collectors (Full Set)	Each defined collector is a tag value. Each object is assigned to the tag value for the collector that it uses when you add the object to vRealize Operations Manager. The default collector is vRealize Operations Manager Collector-vRealize.
Applications (Full Set)	Each defined application is a tag value. When you add a tier to an application, or an object to a tier in an application, the tier is assigned to that tag value.
Maintenance Schedules (Full Set)	Each defined maintenance schedule is a tag value, and objects are assigned to the value when you give them a schedule by adding or editing them.
Adapter Types	Each adapter type is a tag value, and each object that uses that adapter type is given the tag value.
Adapter Instances	Each adapter instance is a tag value, and each object is assigned the tag value for the adapter instance or instances through which its metrics are collected.
Object Types	Each type of object is a tag value, and each object is assigned to the tag value for its type when you add the object.
Recently Added Objects	The last day, seven days, 10 days, and 30 days have tag values. Objects have this tag value as long as the tag value applies to them.
Object Statuses	Tag value assigned to objects that are not receiving data.
Collection States	Tag value assigned to indicate the object collection state, such as collecting or not collecting.
Health Ranges	Good (green), Warning (yellow), Immediate (orange), Critical (red), and Unknown (blue) health statuses have tag values. Each object is assigned the value for its current health status.
Entire Enterprise	The only tag value is Entire Enterprise Applications. This tag value is assigned to each application.

Table 6-1. Predefined Tags (continued)

Tag	Description
Licensing	Tag values are License Groups found under Home > Administration > Management > Licensing. Objects are assigned to the license groups during vRealize Operations Manager installation.
Untag	Drag an object to this tag to delete the tag assignment.

Add an Object Tag and Assign Objects to the Tag

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

Prerequisites

Become familiar with the predefined object tags.

Procedure

- 1 Click **Administration** in the menu, then click **Configuration > Inventory Explorer** in the left pane.
- 2 Click the **Manage Tags** icon above the list of tags.
- 3 Click the **Add New Tag** icon to add a new row and type the name of the tag in the row.
For example, type **Cloud Objects** and click **Update**.
- 4 With the new tag selected, click the **Add New Tag Value** icon to add a new row and type the name of the value in the row.
For example, type **Video Cloud** and click **Update**.
- 5 Click **OK** to add the tag.
- 6 Click the tag to which you want to add objects to display the list of object tag values.
For example, click **Cloud Objects** to display the Video Cloud object tag value.
- 7 Drag objects from the list in the right pane of the Inventory Explorer onto the tag value name.
You can press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.
For example, if you want to assign datacenters that are connected through the vCenter Adapter, type **vCenter** in the search filter and select the datacenter objects to add.

Use a Tag to Find an Object

The quickest way to find an object in vRealize Operations Manager is to use tags. Using tags is more efficient than searching through the entire object list.

Tag values that can also be tags are Applications and Object Types. For example, the Object Types tag has values for each object that is in vRealize Operations Manager, such as Virtual Machine, which includes all the virtual machine objects in your environment. Each of these virtual machines is also a tag value for the Virtual Machine tag. You can expand the tag value list to select the value for which you want to see objects.

Procedure

- 1 In the menu, click **Administration**, then click **Configuration > Inventory Explorer** in the left pane.
- 2 In the tag list in the center pane, click a tag for an object with an assigned value.

When you click a tag, the list of values expands under the tag. The number of objects that is associated with each value appears next to the tag value.

A plus sign next to a tag value indicates that the value is also a tag and that it contains other tag values. You can click the plus sign to see the subvalues.

- 3 Select the tag value.

The objects that have that tag value appear in the pane on the right. If you select multiple tag values, the objects in the list depend on the values that you select.

Tag Value Selection	Objects Displayed
More than one value for the same tag	The list includes objects that have either value. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, the list shows objects that have either value.
Values for two or more different tags	The list includes only objects that have all of the selected values. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, and you also select an adapter instance such as vC-1 of the vCenter Adapter instance tag, only Datacenter or Host System objects associated with vC-1 appear in the list. Datacenter or Host System objects associated with other adapter instances do not appear in the list, nor do objects that are not Datacenter or Host System objects.

- 4 Select the object from the list.

Managing Custom Object Groups in VMware vRealize Operations Manager

A custom object group is a container that includes one or more objects. vRealize Operations Manager uses custom groups to collect data from the objects in the group, and report on the data collected.

Why Use Custom Object Groups?

You use groups to categorize your objects and have the system collect data from the groups of objects and display the results in dashboards and views according to the way you define the data to appear.

You can create static groups of objects, or dynamic groups with criteria that determine group membership as vRealize Operations Manager discovers and collects data from new objects added to the environment.

vRealize Operations Manager provides commonly used object group types, such as World, Environment, and Licensing. The system uses the object group types to categorize groups of objects. You assign a group type to each group so that you can categorize and organize the groups of objects that you create.

Types of Custom Object Groups

When you create custom groups, you can use rules to apply dynamic membership of objects to the group, or you can manually add the objects to the group. When you add an adapter, the groups associated with the adapter become available in vRealize Operations Manager.

- Dynamic group membership. To dynamically update the membership of objects in a group, define rules when you create a group. vRealize Operations Manager adds objects to the group based on the criteria that you define.
- Mixed membership, which includes dynamic and manual.
- Manual group membership. From the inventory of objects, you select objects to add as members to the group.
- Groups associated with adapters. Each adapter manages the membership of the group. For example, the vCenter Server adapter adds groups such as datastore, host, and network, for the container objects in the vSphere inventory. To modify these groups, you must do so in the adapter.

Administrators of vRealize Operations Manager can set advanced permissions on custom groups. Users who have privileges to create groups can create custom groups of objects and have vRealize Operations Manager apply a policy to each group to collect data from the objects and report the results in dashboards and views.

When you create a custom group, and assign a policy to the group, the system uses the criteria defined in the applied policy to collect data from and analyze the objects in the group. vRealize Operations Manager reports on the status, problems, and recommendations for those objects based on the settings in the policy.

Note Only custom groups defined explicitly by users can be exported from or imported to vRealize Operations Manager. Users are able to export or import multiple custom groups. Once an import function has been executed, the user must check to determine if a policy or policies should be associated with the imported group. Export-import operations are available for user defined (created explicitly by user) custom groups only.

How Policies Help vRealize Operations Manager Report On Object Groups

When you apply a policy to an object group, vRealize Operations Manager uses threshold settings, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that you enabled in the policy to collect data from the objects in the group, and report the results in dashboards and views.

When you create a new object group, you have the option to apply a policy to the group.

- To associate a policy with the custom object group, select the policy in the group creation wizard.
- To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

User Scenario: Creating Custom Object Groups

As a system administrator, you must monitor the capacity for your clusters, hosts, and virtual machines. vRealize Operations Manager monitors them at different service levels to ensure that these objects adhere to the policies established for your IT department, and discovers and monitors new objects added to the environment. You have vRealize Operations Manager apply policies to the object groups to analyze, monitor, and report on the status of their capacity levels.

To have vRealize Operations Manager monitor the capacity levels for your objects to ensure that they adhere to your policies for your service levels, you categorize your objects into Platinum, Gold, and Silver object groups to support the service tiers established.

You create a group type, and create dynamic object groups for each service level. You define membership criteria for each dynamic object group to have vRealize Operations Manager keep the membership of objects current. For each dynamic object group, you assign the group type, and add criteria to maintain membership of your objects in the group. To associate a policy with the custom object group, you can select the policy in the group creation wizard.

Prerequisites

- Know the objects that exist in your environment, and the service levels that they support.
- Understand the policies required to monitor your objects.
- Verify that policies are available to monitor the capacity of your objects.

Procedure

- 1 To create a group type to identify service level monitoring, click **Administration** in the menu, then click **Configuration > Group Types**.

- 2 On the Group Types toolbar, click the plus sign and type **Service Level Capacity** for the group type.

Your group type appears in the list.

- 3 Click **Environment** in the menu, then click the **Custom Groups** tab.

- 4 To create a new object group, click the **plus** sign on the Groups toolbar.

The New Group workspace appears where you define the data and membership criteria for the dynamic group.

- a In the Name text box, type a meaningful name for the object group, such as **Platinum_Objects**.
- b In the **Group Type** drop-down menu, select **Service Level Capacity**.
- c (Optional) In the **Policy** drop-down menu, select your service level policy that has thresholds set to monitor the capacity of your objects.

To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

- d Select the **Keep group membership up to date** check box so that vRealize Operations Manager can discover objects that meet the criteria, and add those objects to the group.
- 5 Define the membership for virtual machines in your new dynamic object group to monitor them as platinum objects.
 - a From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Virtual Machine**.
 - b From the empty drop-down menu for the criteria, select **Metrics**.
 - c From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
 - d From the conditional value drop-down menu, select **is less than**.
 - e From the **Metric value** drop-down menu, type **10**.
 - 6 Define the membership for host systems in your new dynamic object group to monitor them as platinum objects.
 - a Click **Add another criteria set**.
 - b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Host System**.
 - c From the empty drop-down menu for the criteria, select **Metrics**.

- d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
 - e From the conditional value drop-down menu, select **is less than**.
 - f From the **Metric value** drop-down menu, type **100**.
- 7** Define the membership for cluster compute resources in your new dynamic object group.
- a Click **Add another criteria set**.
 - b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Cluster Compute Resources**.
 - c From the empty drop-down menu for the criteria, select **Metrics**.
 - d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **capacityRemaining**.
 - e From the conditional value drop-down menu, select **is less than**.
 - f From the **Metric value** drop-down menu, type **1000**.
 - g Click **Preview** to determine whether objects already match this criteria.
- 8** Click **OK** to save your group.

When you save your new dynamic group, the group appears in the Service Level Capacity folder, and in the list of groups on the **Groups** tab.

- 9** Wait five minutes for vRealize Operations Manager to collect data from the objects in your environment.

Results

vRealize Operations Manager collects data from the cluster compute resources, host systems, and virtual machines in your environment, according to the metrics that you defined in the group and the thresholds defined in the policy that is applied to the group, and displays the results about your objects in dashboards and views.

What to do next

To monitor the capacity levels for your platinum objects, create a dashboard, and add widgets to the dashboard. See [Dashboards](#).

Managing Application Groups

An application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. vRealize Operations Manager builds an application to determine how your environment is affected when one or more components in an application experiences problems, and to monitor the overall health and performance of the application. Object membership in an application is not dynamic. To change the application, you manually modify the objects in the container.

Reasons to Use Applications

vRealize Operations Manager collects data from components in the application and displays the results in a summary dashboard for each application with a real-time analysis for any of the components. If a component experiences problems, you can see where in the application the problems arise, and determine how problems spread to other objects.

Note vRealize Operations Manager provides for calendar periodicity. If your application includes work performed on a specific day of the month, for example, the 15th of the month or the last day of the month, this calendar function identifies the pattern after six cycles of the application. Once the pattern is recognized, the system can forecast accurately into the future. Because the system acquires its information from the input data, you do not have to give any details about how you schedule periodical work.

User Scenario: Adding an Application

As the system administrator of an online training system, you must monitor components in the Web, application, and database tiers of your environment that can affect the performance of the system. You build an application that groups related objects together in each tier. If a problem occurs with one of the objects, it is reflected in the application display and you can open a summary to investigate the source of the problem further.

In your application, you add the DB-related objects that store data for the training system in a tier, Web-related objects that run the user interface in a tier, and application-related objects that process the data for the training system in a tier. The network tier might not be needed. Use this model to develop your application.

Procedure

- 1 In the menu, click **Environment**, then click **Groups and Applications** in the left pane.
- 2 Click the **Applications** tab and click the plus sign.
- 3 Click **Basic n-tier Web App** and click **OK**.

The Application Management page that appears has two rows. Select objects from the bottom row to populate the tiers in the top row.

- 4 Type a meaningful name such as **Online Training Application** in the Application text box.
- 5 For each of the Web, application and database tiers listed, add the objects to the Tier Objects section.
 - a Select a tier name. This is the tier that you populate.
 - b To the left of the object row, select object tags to filter for objects that have that tag value. Click the tag name once to select the tag from the list and click the tag name again to deselect the tag from the list. If you select multiple tags, objects displayed depend on the values that you select.

You can also search for the object by name.

- c To the right of the object row, select the objects to add to the tier.
 - d Drag the objects to the Tier Objects section.
- 6** Click Save to save the application.

Results

The new application appears in the list of applications on the Environment Overview Applications page. If any of the components in any of the tiers develops a problem, the application displays a yellow or red status.

What to do next

To investigate the source of the problem, click the application name and evaluate the object summary information. See the *vRealize Operations Manager User Guide*.

Configuring Data Display

7

You configure the content in vRealize Operations Manager to suit your information needs, using views, reports, dashboards, and widgets.

Views display data, based on an object type. You can select from various view types to see your data from a different perspective. Views are reusable components that you can include in reports and dashboards. Reports can contain predefined or custom views and dashboards in a specified order. You build the reports to represent objects and metrics in your environment. You can customize the report layout by adding a cover page, a table of contents, and a footer. You can export the report in a PDF or CSV file format for further reference.

You use dashboards to monitor the performance and state of objects in your virtual infrastructure. Widgets are the building blocks of dashboards and display data about configured attributes, resources, applications, or the overall processes in your environment. You can also incorporate views in dashboards using the vRealize Operations Manager View Widget.

This chapter includes the following topics:

- [Widgets](#)
- [Dashboards](#)
- [Views](#)
- [Reports](#)

Widgets

Widgets are the panes on your dashboards. You add widgets to a dashboard to create a dashboard. Widgets show information about attributes, resources, applications, or the overall processes in your environment.

You can configure widgets to reflect your specific needs. The available configuration options vary depending on the widget type. You must configure some of the widgets before they display any data. Many widgets can provide or accept data from one or more widgets. You can use this feature to set the data from one widget as filter and display related information on a single dashboard.

Widget Interactions

Widget interactions are the configured relationships between widgets in a dashboard where one widget provides information to a receiving widget. When you are using a widget in the dashboard, you select data on one widget to limit the data that appears in another widget, allowing you to focus on a smaller subset data.

How Interactions Work

If you configured interactions between widget at the dashboard level, you can then select one or more objects in the providing widget to filter the data that appears in the receiving widget, allowing you to focus on data related to an object.

To use the interaction option between the widgets in a dashboard, you configure interactions at the dashboard level. If you do not configure any interactions, the data that appears in the widgets is based on how the widget is configured.

When you configure widget interaction, you specify the providing widget for the receiving widget. For some widgets, you can define two providing widgets, each of which can be used to filter data in the receiving widget.

For example, if you configured the Object List widget to be a provider widget for the Top-N widget, you can select one or more objects in the Object List widget and the Top-N displays data only for the selected objects.

For some widgets, you can define more than one providing widget. For example, you can configure the Metric Chart widget to receive data from a metrics provider widget and an objects providing widget. In such case, the Metric Chart widget shows data for any object that you select in the two provider widgets.

Manage Metric Configuration

You can create a custom set of metrics to display the widgets. You can configure one or more files that define different sets of metrics for a particular adapter and object types so that the supported widgets are populated based on the configured metrics and selected object type.

How the Metric Configuration Works

From the Metric Configuration page, you create an XML file that displays a set of metrics at a supported widget. The widgets are Metric Chart, Property List, Rolling View Chart, Scoreboard, Sparkline Chart, and Topology Graph. To use the metric configuration, you must set the widget Self Provider to **Off** and create a widget interaction with a provider widget.

Where You Find the Metric Configuration

To manage metric configurations, in the menu, click **Administration**, and then in the left pane click **Configuration > Metric Configurations**.

Table 7-1. Manage Metric Config Toolbar Options

Option	Description
Create Configuration	Creates an empty XML file in a selected folder.
Edit Configuration	Activates a selected XML file for edit in the text box on the right.
Delete Configuration	Deletes a selected XML file.
Text box	Displays a selected XML file. You must select an XML file and click Edit to edit it.

Add a Resource Interaction XML File

A resource interaction file is a custom set of metrics that you want to display in widgets that support the option. You can configure one or more files that define different sets of metrics for particular object types so that the supported widgets are populated based the configured metrics and selected object type.

The following widgets support the resource interaction mode:

- Metric Chart
- Property List
- Rolling View Chart
- Scoreboard
- Sparkline Chart
- Topology Graph

To use the metric configuration, which displays a set of metrics that you defined in an XML file, the dashboard and widget configuration must meet the following criteria:

- The dashboard **Widget Interaction** options are configured so that another widget provides objects to the target widget. For example, an Object List widget provides the object interaction to a chart widget.
- The widget **Self Provider** option is set to **Off**.
- The custom XML file in the **Metric Configuration** drop-down menu is in the `/usr/lib/vmware-vcops/tools/opsccli` directory and has been imported into the global storage using the `import` command.

If you add an XML file and later modify it, the changes might not take effect.

Prerequisites

- Verify that you have the necessary permissions to access the installed files for vRealize Operations Manager and add files.

- Create a new files based on the existing examples. Examples are available in the following location:
 - vApp. The XML file is in `/usr/lib/vmware-vcops/tomcat-web-app/webapps/vcops-web-ent/WEB-INF/classes/resources/reskndmetrics`.

Procedure

- 1 Create an XML file that defines the set of metrics.

For example:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<AdapterKinds>
  <AdapterKind adapterKindKey="VMWARE">
    <ResourceKind resourceKindKey="HostSystem">
      <Metric attrkey="sys:host/vim/vmvisor/slp|resourceMemOverhead_latest" />
      <Metric attrkey="cpu|capacity_provisioned" />
      <Metric attrkey="mem|host_contention" />
    </ResourceKind>
  </AdapterKind>
</AdapterKinds>
```

In this example, the displayed data for the host system based on the specified metrics.

- 2 Save the XML file in one of the following directories base on the operating system of your vRealize Operations Manager instance.

Operating System	File Location
vApp	<code>/usr/lib/vmware-vcops/tools/opscli</code>

- 3 Run the import command.

Operating System	File Location
vApp	<code>./ops-cli.sh file import reskndmetric YourCustomFilename.xml</code>

The file is imported into global storage and is accessible from the supported widgets.

- 4 If you update an existing file and must re-import the file, append `--force` to the above import command and run it.

For example, `./vcops-cli.sh file import reskndmetric YourCustomFilename.xml --force`.

What to do next

To verify that the XML file is imported, configure one of the supported widgets and ensure that the new file appears in the drop-down menu.

You can also create a custom set of metrics to display the widgets, from the [Manage Metric Configuration](#).

Widget Definitions List

A widget is a pane on a dashboard that contains information about configured attributes, resources, applications, or the overall processes in your environment. Widgets can provide a holistic, end-to-end view of the health of all of the objects and applications in your enterprise. If your user account has the necessary access rights, you can add and remove widgets from your dashboards.

Table 7-2. Summary of Widgets

Widget Name	Description
Alert List	Shows a list of alerts for the objects that the widget is configured to monitor. If no objects are configured, the list displays all alerts in your environment.
Alert Volume	Shows a trend report for the last seven days of alerts generated for the objects it is configured to monitor.
Anomalies	Shows a chart of the anomalies count for the past 6 hours.
Anomaly Breakdown	Shows the likely root causes for symptoms for a selected resource.
Capacity Remaining	Shows a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.
Container Details	Shows the health and alert counts for each tier in a single selected container.
Container Overview	Shows the overall health and the health of each tier for one or more containers.
Current Policy	Shows the highest priority policy applied to a custom group.
Data Collection Results	Shows a list of all supported actions specific for a selected object.
DRS Cluster Settings	Shows the workload of the available clusters and the associated hosts.
Efficiency	Shows the status of the efficiency-related alerts for the objects that it is configured to monitor. Efficiency is based on generated efficiency alerts in your environment.
Environment	Lists the number of resources by object or groups them by object type.
Environment Overview	Shows the performance status of objects in your virtual environment and their relationships. You can click an object to highlight its related objects and double-click an object to view its Resource Detail page.
Environment Status	Shows statistics for the overall monitored environment.
Faults	Shows a list of availability and configuration issues for a selected resource.
Forensics	Shows how often a metric had a particular value, as a percentage of all values, within a given time period. It can also compare percentages for two time periods.
Geo	Shows where your objects are located on a world map, if your configuration assigns values to the Geo Location object tag.
Health	Shows the status of the health-related alerts for the objects that it is configured to monitor. Health is based on generated health alerts in your environment.
Health Chart	Shows health information for selected resources, or all resources that have a selected tag.
Heat Map	Shows a heat map with the performance information for a selected resource.
Mashup Chart	Brings together disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs). This widget is typically used for a container.

Table 7-2. Summary of Widgets (continued)

Widget Name	Description
Metric Chart	Shows a chart with the workload of the object over time based on the selected metrics.
Metric Picker	Shows a list of available metrics for a selected resource. It works with any widget that can provide resource ID.
Object List	Shows a list of all defined resources.
Object Relationship	Shows the hierarchy tree for the selected object.
Object Relationship (Advanced)	Shows the hierarchy tree for the selected objects. It provides advanced configuration options.
Property List	Shows the properties and their values of an object that you select.
Recommended Actions	Displays recommendations to solve problems in your vCenter Server instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines.
Risk	Shows the status of the risk-related alerts for the objects that it is configured to monitor. Risk is based on generated risk alerts in your environment.
Rolling View Chart	Cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget.
Scoreboard	Shows values for selected metrics, which are typically KPIs, with color coding for defined value ranges.
Scoreboard Health	Shows color-coded health, risk, and efficiency scores for selected resources.
Sparkline Chart	Shows graphs that contain metrics for an object . If all of the metrics in the Sparkline Chart widget are for an object that another widget provides, the object name appears at the top right of the widget.
Tag Picker	Lists all defined resource tags.
Text Display	Reads text from a Web page or text file and shows the text in the user interface.
Time Remaining	Shows a chart of the Time Remaining values for a specific resources over the past 7 days.
Top Alerts	Lists the alerts most likely to negatively affect your environment based on the configured alert type and objects.
Top-N	Shows the top or bottom N number metrics or resources in various categories, such as the five applications that have the best or worst health.
Topology Graph	Shows multiple levels of resources between nodes.
View	Shows a defined view depending on the configured resource.
Weather Map	Uses changing colors to show the behavior of a selected metric over time for multiple resources.
Workload	Shows workload information for a selected resource.
Workload Pattern	Shows a historical view of the hourly workload pattern of an object.
Workload Utilization	Shows the workload utilization for objects so that you can identify problems with workload.

For more information about the widgets, see the vRealize Operations Manager help.

Dashboards

Dashboards present a visual overview of the performance and state of objects in your virtual infrastructure. You use dashboards to determine the nature and timeframe of existing and potential issues with your environment. You create dashboards by adding widgets to a dashboard and configuring them.

vRealize Operations Manager collects performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems. The data and analysis are presented through alerts, in configurable dashboards, on predefined pages, and in several predefined dashboards.

- You can start with several predefined dashboards in vRealize Operations Manager.
- You can create extra ones that meet your specific needs using widgets, views, badges, and filters to change the focus of the information.
- You can clone and edit the predefined dashboards or start from scratch.
- To display data that shows dependencies, you can add widget interactions in dashboards.
- You can provide role-based access to various dashboards for better collaboration in teams.

Table 7-3. Menu Options

Menu	Description
All Dashboards	Lists the dashboards that are enabled. You can use this menu for a quick navigation through your dashboards. When you navigate to a dashboard using the All Dashboards option, the dashboard is listed in the left pane of the Dashboards page.
Actions	Available dashboard actions, such as create, edit, delete, and set as default. These actions are applied directly to the dashboard that you are on.

Types Of Dashboards

You can use the predefined dashboards or create your own custom dashboard in vRealize Operations Manager.

Custom Dashboards

vRealize Operations Manager has predefined dashboards. You can also create dashboards that meet your environment needs.

To manage your dashboards, in the menu, click **Dashboards**. Click **Actions > Manage Dashboards** and then click the gear icon.

Depending on your access rights, you can add, delete, and arrange widgets on your dashboards. You can also clone and create dashboards, import or export dashboards from other instances, edit widget configuration options, configure widget interactions, and transfer ownership of dashboards.

Table 7-4. Dashboards Options

Option	Description	Usage
Save as Template	Contains all the information in a dashboard definition.	You can use any dashboard to create a template.
Export Dashboards	When you export a dashboard, vRealize Operations Manager creates a dashboard file in JSON format.	You can export a dashboard from one vRealize Operations Manager instance and import it to another.
Import Dashboards	A PAK or JSON file that contains dashboard information from vRealize Operations Manager.	You can import a dashboard that was exported from another vRealize Operations Manager instance.
Enable Dashboard(s)	Enables a dashboard that was previously disabled.	
Disable Dashboard(s)	Disables a dashboard.	
Transfer Dashboard(s)	Assigns a new owner to the dashboard.	After you assign a dashboard to a new owner, the dashboard is no longer displayed as one of your dashboards. When you transfer a dashboard that was previously shared with user groups, information about the shared user groups and group hierarchy is retained.
Remove Dashboard(s) from Home	Removes a dashboard from the vRealize Operations Manager home page.	You can add any dashboard to the vRealize Operations Manager home page.
Reorder/Autoswitch Dashboards	Changes the order of the dashboard tabs on vRealize Operations Manager home page.	You can configure vRealize Operations Manager to switch from one dashboard to another.
Manage Summary Dashboards	Provides you with an overview of the state of the selected object, group, or application.	You can change the Summary tab with a dashboard to get information specific to your needs.
Manage Dashboard Groups	Groups dashboards in folders.	You can create dashboard folders to group the dashboards in a way that is meaningful to you.
Share Dashboards	Makes a dashboard available to other users or user groups.	You can share a dashboard or dashboard template with one or more user groups.
Copy Dashboards	Copies a dashboard to other another user or user group.	You can copy a dashboard to another user or user group. Specify the dashboards to be shared and select a target user and specify the target folder.

The dashboard list depends on your access rights.

Predefined Dashboards

vRealize Operations Manager has predefined dashboards that address several key questions including how you can troubleshoot your VMs, the workload distribution of your hosts, clusters,

and datastores, the capacity of your data center, and information about the VMs. You can also view log details.

In the menu, click **Dashboards**, and then in the left pane you can access a list of predefined dashboards. For dashboards to appear in the left pane of the Dashboards page, in the menu, click **Dashboards**, and then from the **All Dashboards** drop-down menu select the check-box against the required dashboard.

The default dashboard that appears when you click **Dashboards** in the menu is the **Getting Started** dashboard. You can close a dashboard from the left pane by selecting the dashboard and clicking the **X** icon. The dashboard you last opened is displayed the next time you navigate to **Dashboards** in the menu. If there is only one dashboard left in the left pane, you cannot close it.

The following predefined dashboards can be accessed by clicking **Dashboards** in the menu, and then clicking **All Dashboards**:

- Capacity and Utilization
 - Capacity Allocation Overview
 - Cluster Utilization
 - Datastore Utilization
 - Heavy Hitter VMs
 - Host Utilization
 - Utilization Overview
 - VM Utilization
 - vSAN Capacity Overview
- Configuration and Compliance
 - Cluster Configuration
 - Distributed Switch Configuration
 - Host Configuration
 - VM Configuration
 - vSphere Hardening Compliance
- Operations
 - Datastore Usage Overview
 - Host Usage Overview
 - Migrate to vSAN
 - Operations Overview
 - vSAN Operations Overview

- Optimize
 - Assess Cost
 - Optimization History
 - Optimize Performance
- Performance Troubleshooting
 - Troubleshoot a Cluster
 - Troubleshoot a Datastore
 - Troubleshoot a Host
 - Troubleshoot a VM
 - Troubleshoot vSAN
 - Troubleshoot with Logs
- vRealize Assessments
 - Hybrid Cloud Assessment
 - vSphere Optimization Assessment
- vRealize Automation
 - Application Overview
 - Environment Overview
 - Resource Consumption Overview
 - Top-N
- vRealize Operations
 - MP Statistics
 - Self Cluster Statistics
 - Self Health
 - Self Performance Details
 - Self Services Communications
 - Self Services Summary
 - Self Troubleshooting
 - vCenter Adapter Details
- Getting Started

Getting Started Dashboard

The Getting Started dashboard is a guide to answering the most frequent questions of your IT staff. The dashboard breaks tasks into broad categories including Capacity and Utilization, Configuration and Compliance, Operations, Performance Troubleshooting, and Optimize. Using each of these categories you can drill down to the specific use cases and problems you are trying to solve. Each problem statement is associated with a predefined dashboard that you can access through this page. To view a dashboard, click the dashboard name listed on the right side of the Getting Started dashboard.

Capacity and Utilization Dashboards

The dashboards in the Capacity and Utilization category cater to the teams responsible for tracking the utilization of the provisioned capacity in their virtual infrastructure. The dashboards within this category allow you to take capacity procurement decisions, reduce wastage through reclamation, and track usage trends to avoid performance problems due to capacity shortfalls.

Key questions these dashboards help you answer are as follows:

- How much capacity exists, how much is used, and the usage trends for a specific vCenter, data center, or cluster?
- How much disk, vCPU, or memory you can reclaim from large VMs in your environment to reduce wastage and improve performance?
- Which clusters have the highest resource demands?
- Which hosts are being heavily utilized and why?
- Which datastores are running out of disk space and who are the top consumers?
- The storage capacity and utilization of your vSAN environment with the savings achieved by enabling deduplication and compression.

Capacity Allocation Overview Dashboard

This dashboard provides an overview of allocation ratios for virtual machines, vCPUs, and memory for a specific data center or cluster.

Cluster Utilization Dashboard

The Cluster Utilization dashboard helps you identify vSphere clusters that are extensively consumed from a CPU, memory, disk, and network perspective.

You can use this dashboard to identify the clusters that cannot serve the virtual machine demand.

You can select a cluster with high CPU, memory, disk, or network demand. The dashboard lists the ESXi hosts that are a part of the given cluster. If there is an imbalance in the use of hosts within the selected clusters, you can balance the hosts by moving the VMs within the cluster.

You can use this dashboard to view the historical cluster demand. If the situation is critical, use Workload Balance and move the VMs out of the clusters to avoid potential performance issues. For more information, see [Chapter 3 Configuring and Using Workload Optimization](#). If all the clusters in a given environment display the same pattern, you might have to add new capacity to cater to the increase in demand.

Datastore Utilization Dashboard

The Datastore Utilization dashboard helps you identify storage provisioning and utilization patterns in a virtual infrastructure.

As a best practice, ensure that the datastores are of standard size, to manage storage in your virtual environments. The heat map on this dashboard displays all the datastores monitored by vRealize Operations Manager and groups them by clusters.

The dashboard uses colors to depict the utilization pattern of the datastores. Grey represents an underutilized datastore, red represents a datastore that has run out of disk space, and green represents an optimally used datastore. You can select a datastore from the dashboard to see the past utilization trends and forecasted usage. The dashboard lists all the VMs that run on the selected datastore. You can reclaim storage used by large VM snapshots or powered off VMs.

You can use the vRealize Operations Manager action framework to reclaim resources by deleting the snapshots or unwanted powered off VMs.

- **Datastore Capacity and Utilization:** Use this widget to find out which datastores are overused and which ones are underused. You can also find out whether the datastores are of equal size. When you select a datastore from this widget, the dashboard is automatically populated with the relevant data.
- **VMs in the Selected Datastore:** Use this widget to view a list of VMs based on the datastore you select. You can also view relevant details such as whether the VMs are powered on and the size of the snapshot if any.
- **Usage Trend of Selected Datastore:** Use this widget to find out the trends in capacity used by a selected datastore as against the total capacity available.
- **All Shared Datastores in the Environment:** Use this widget to view a list of datastores that are shared in your environment. The information displayed in this widget helps you make an informed decision about whether you have to rebalance the capacity of the datastores based on usage.

Heavy Hitter VMs

The Heavy Hitter VMs dashboard helps you identify virtual machines which are consistently consuming a large amount of resources from your virtual infrastructure. In heavily over-provisioned environments, this might create resource bottlenecks resulting in potential performance issues.

You can use this dashboard to identify the resource utilization trends of each of your vSphere clusters. With the utilization trends, you can also view a list of VMs within those clusters based on their resource demands from the CPU, memory, disk, and network within your environment. You can also analyze the workload pattern of these VMs over the past week to identify heavy hitter VMs which might be running a sustained, heavy workload that is measured over a day, or bursty workloads that is measured using peak demand.

You can export a list of offenders and take appropriate action to distribute this demand and reduce potential bottlenecks.

You can use the dashboard widgets in several ways.

- **Select a Cluster:** Use this widget to select a cluster. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Cluster CPU and Cluster Memory:** Use these widgets to view the CPU and memory for the cluster.
- **Cluster IOPS and Cluster Network Throughput:** Use these widgets to view the IOPS and network throughput for the cluster.
- Use the other widgets in the dashboard to view which VMs in the cluster generated the highest network throughput and IOPS. You can also view which VMs in the cluster generated the highest CPU demand and the highest memory demand. You can compare the information for the VM with the results for the cluster and correlate the trends. You can manually set the time to the time period for which you want to view data.

Host Utilization Dashboard

The Host Utilization dashboard helps you identify hosts that are extensively consumed from a CPU, memory, disk, and network perspective.

You can use this dashboard to identify hosts that cannot serve the virtual machine demand. The dashboard provides a list of the top 10 virtual machines. You can identify the source of this unexpected demand and take appropriate actions.

You can use the dashboard to view demand patterns over the last 24 hours and identify hosts that have a history of high demand. You must move the virtual machines out of these hosts to avoid potential performance issues. If all the hosts of a given cluster display the same pattern, you might have to add new capacity to cater to the increase in demand.

Utilization Overview Dashboard

The Utilization Overview dashboard helps you view the available capacity in the virtual infrastructure.

The Utilization Overview dashboard allows you to assess the utilization at each resource group level such as vCenter, data center, custom data center, or vSphere cluster. You can quickly select an object and view the total capacity, used capacity, and usable capacity of the object to understand the current capacity situation.

You can use the dashboard widgets in several ways.

- **Total Environment Summary:** Use this widget to view the total available capacity in the environment including information about the number of hosts and datastores. You can also view storage, memory, and CPU capacity, and the number of physical CPUs.
- **Select an Environment:** Use this widget to select a data center, a cluster compute resource, or a vCenter Server. You can use the filter to narrow your list based on several parameters. After you identify the data center you want to view, select it. The dashboard is populated with the relevant data.
- **Inventory:** Use this widget to view the number of running VMs and hosts. You can also view the number of datastores and the consolidation ratio in the environment.

- **Usable Capacity (Exclude HA Buffers):** Use this widget to view the capacity that is available in the virtual infrastructure.
- **Used Capacity:** Used this widget to view how the capacity is used in various data centers and clusters.
- **Capacity Remaining:** Use this widget to view the capacity remaining in terms of memory, storage, and CPU capacity remaining.
- **Predicted Time Remaining:** Use this widget to view the predicted time remaining based on the use patterns in the environment.
- **Cluster Capacity Details:** Use this widget to view detailed capacity information for each cluster.

VM Utilization Dashboard

The VM Utilization dashboard helps you as an administrator to capture the utilization trends of any VM in your environment. You can list the key properties of a VM and the resource utilization trends for a specific time period. You can share the details with the VM or application owners.

The dashboard displays resource utilization trends so that the VM or application owners can view these trends when they expect a high load on applications. For example, activities like batch jobs, backup schedules, and load testing. Application owners must ensure that the VMs do not consume 100% of the provisioned resources during these periods. Excessive consumption of the provisioned resources can lead to resource contention within the applications and can cause performance issues.

- **Search for a VM to Report its Usage:** Use this widget to select the VM you want to troubleshoot. You can use the filter to narrow your list based on several parameters. After you identify the VM that you want to view, select it. The dashboard is automatically populated with the relevant data.
- **About the VM:** Use this widget to view the VM you selected and its details. You select the VM in the Search for a VM to Report its Usage widget.
- **VM Utilization Trend: CPU, Memory, IOPS, Network:** Use this widget to view information about the utilization and allocation trends for CPU demand, memory workload, disk commands per second, and the network usage rate.

vSAN Capacity Overview

The vSAN Capacity Overview dashboard provides an overview of vSAN storage capacity and savings achieved by enabling deduplication and compression across all vSAN clusters.

You can view current and historical use trends, and future procurement requirements from the dashboard. You can view details such as capacity remaining, time remaining, and storage reclamation opportunities to make effective capacity management decisions.

You can view the distribution of use among vSAN disks from the dashboard. You can view these details either as an aggregate or at an individual cluster level.

Configuration and Compliance Dashboards

The dashboards in the Configuration and Compliance category cater to administrators who are responsible for managing configuration drifts within a virtual infrastructure. Since most of the issues in a virtual infrastructure are a result of inconsistent configurations, dashboards in this category highlight the inconsistencies at various levels such as VMs, hosts, clusters, and virtual networks. You can view a list of configuration improvements that helps you avoid problems that are caused because of misconfigurations.

Your IT security teams can also measure your environment against the vSphere hardening best practices to ensure that your environment is fully secured and meets all the compliance standards.

Key questions these dashboards help you answer are as follows:

- Are the vSphere clusters consistently configured for high availability (HA) and optimal performance?
- Are the ESXi hosts consistently configured and available to use?
- Are the VMs sized and configured as per the recommended best practices?
- Are virtual switches configured optimally?
- Is the environment configured in accordance with the vSphere Hardening Guide?

Cluster Configuration Dashboard

The Cluster Configuration dashboard provides a quick overview of your vSphere cluster configurations. The dashboard highlights the areas that are important in delivering performance and availability to your virtual machines. The dashboard also highlights if there are clusters which are not configured for DRS, High Availability (HA), or admission control to avoid any resource bottlenecks or availability issues when a host fails.

The heat map in this dashboard helps you to identify if you have hosts where vMotion was not enabled as this may not allow the VMs to move from or to that host. This may cause potential performance issues for the VMs on that host if the host gets too busy. You can also view how consistently your clusters are sized and whether the hosts on each of those clusters are consistently configured.

The Cluster Properties widget in this dashboard allows you to report on all these parameters by exporting the data. You can share the data with the relevant stakeholders within your organization.

You can use the dashboard widgets in several ways.

- **vSphere DRS Status, vSphere HA Status, and HA Admission Control Status:** Use these widgets to view if there are clusters that are not configured for DRS, HA, or admission control. With the information, you can avoid resource bottlenecks or availability issues when a host fails.
- **Is vMotion enabled on hosts in a cluster:** Use this widget to identify if you have hosts where vMotion was not enabled. If vMotion is not enabled, the VMs do not move from or to the host and causes potential performance issues in the VMs on that host if the host gets too busy.

- **Host Count across Clusters:** Use this widget to view all the clusters in your environment. If the clusters have a consistent number of hosts, the boxes displayed are of equal size. This representation helps you determine whether there is a large deviation among cluster sizes, whether there is a small cluster with fewer than four hosts, or whether there is a large cluster. Operationally, keep your clusters consistent and of moderate size.
- **Attributes of ESXi Hosts in the Selected Cluster:** Use this widget to view the configuration details for the hosts within a cluster.
- **All Clusters Properties:** Use this widget to view the properties for all the clusters in the widget.

Distributed Switch Configuration Dashboard

The Distributed Switch Configuration dashboard allows you to view details of virtual switch configuration and utilization. When you select a virtual switch, you can see the list of ESXi hosts, distributed port groups, and virtual machines that use or are on the selected switch. You can also find out which ESXi hosts and VMs use a specific switch.

You can identify misconfigurations within various network components by reviewing the properties listed in the views within the dashboard. You can track important information such as the IP address and the MAC address assigned to the virtual machines.

As a network administrator, you can use this dashboard to get visibility into the virtual infrastructure network configuration.

You can use the dashboard widgets in several ways.

- **Select a Distributed Switch:** Use this widget to select the switch for which you want to view details. You can use the filter to narrow your list based on several parameters. After you identify the switch that you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Distributed Port Groups on the Switch:** Use this widget to view the port groups on the switch, how many ports each switch has, and the usage details.
- **ESXi Hosts/VMs Using the Selected Switch:** Use these widgets to find out which ESXi hosts and VMs use the selected switch. You can also view configuration details about the ESXi hosts and VMs that use the selected switch.

Host Configuration Dashboard

The Host Configuration dashboard provides an overview of your ESXi host configurations, and displays inconsistencies so that you can take corrective action.

The dashboard also measures the ESXi hosts against the vSphere best practices and indicates deviations that can impact the performance or availability of your virtual infrastructure. Although you can view this type of data in other dashboards, in this dashboard you can export the ESXi configuration view and share it with other administrators.

VM Configuration Dashboard

The VM dashboard focuses on highlighting the key configurations of the virtual machines in your environment. You can use this dashboard to find inconsistencies in configuration within your

virtual machines and take quick remedial measures. You can safeguard the applications which are hosted on these virtual machines by avoiding potential issues due to misconfigurations.

Some of the basic problems the dashboard focuses on includes identifying VMs running on older VMware tools versions, VMware tools not running, or virtual machines running on large disk snapshots. VMs with such symptoms can lead to potential performance issues and hence it is important that you ensure that they do not deviate from the defined standards. This dashboard includes a predefined Virtual Machine Inventory Summary report which you can use to report the configurations highlighted in this dashboard for quick remediation.

You can use the dashboard widgets in several ways.

- Use the Large VMs widgets to view graphical representations of VMs that have a large CPU, RAM, and disk space.
- **Guest OS Distribution:** Use this widget to view a break up of the different flavors of operating systems you are running.
- **Guest Tools Version** and **Guest Tools Status:** Use these widgets to identify if you have inconsistent or older version of VMware tools which might lead to performance issues.
- View the VMs with limits, large snapshots, orphaned VMs, VMs with more than one NIC, and VMs with a nonstandard operating system. These VMs have a performance impact on the rest of the VMs in your environment even though they do not fully use their allocated resources.

You can customize the views in the widgets.

- 1 Click the **Edit Widget** icon from title bar of the widget. The **Edit** widget dialog box is displayed.
- 2 From the **Views** section, click the **Edit View** icon. The **Edit View** dialog box is displayed.
- 3 Click the **Presentation** option in the left pane and make the required modifications.

vSphere Security Compliance Dashboard

The vSphere Security Compliance dashboard measures your environment against the *vSphere Hardening Guide* and lists any objects which are non-compliant.

This dashboard displays the trend of high risk, medium risk, and low risk violations and shows the overall compliance score of your virtual infrastructure. Using heat maps, you can investigate various components to check the compliance for your ESXi hosts, clusters, port groups, and virtual machines. Each non-compliant object is listed in the dashboard with recommendations on the remediation required to secure your environment.

Operations Dashboards

The dashboards in the Operations category are most helpful to personnel within an organization that require a summary of important data to take quick decisions. As a member of the network operations center (NOC) team, you may want to identify problems and take action or as an executive, you may want a quick overview of your environments to keep track of important KPIs.

Key questions these dashboards help you answer are as follows:

- What does the infrastructure inventory look like?
- What is the alert volume trend in the environment?
- Are virtual machines being served well?
- Are there areas in the data center you have to worry about?
- What does the vSAN environment look like and are there optimization opportunities by migrating VMs to vSAN?

Datastore Usage Overview Dashboard

The Datastore Usage Overview dashboard provides a view of all the virtual machines in your environment in a heat map. The dashboard is suitable for an NOC environment.

The heat map contains a box for each virtual machine in your environment. You can identify the virtual machines that are generating excessive IOPS because the boxes are sized by the number of IOPS they generate.

The colors of the boxes represent the latency experienced by the virtual machines from the underlying storage. An NOC administrator can investigate the cause of this latency and resolve it to avoid potential performance problems.

Host Usage Overview Dashboard

The Host Usage Overview dashboard provides a view of all the ESXi hosts in your environment in a heat map. The dashboard is suitable for an NOC environment.

Using this dashboard an NOC administrator can easily find resource bottlenecks created due to excessive Memory Demand, Memory Consumption or CPU Demand.

The heat map displays hosts grouped by clusters to help you locate clusters that are using excessive CPU or memory. You can also identify if you have ESXi hosts within the clusters that are not evenly utilized. An administrator can then trigger activities such as workload balance or set DRS to ensure that hot spots are eliminated.

Migrate to vSAN

The Migrate to vSAN dashboard provides you with an easy way to move virtual machines from existing storage to newly deployed vSAN storage.

You can use this dashboard to select non-vSAN datastores that might not serve the virtual machine IO demand. By selecting the virtual machines on a given datastore, you can identify the historical IO demand and the latency trends of a given virtual machine. You can then find a suitable vSAN datastore which has the space and the performance characteristics to serve the demand of this VM. You can move the virtual machine from the existing non-vSAN datastore to the vSAN datastore. You can continue to watch the use patterns to see how the VM is served by vSAN after you move the VM.

Operations Overview Dashboard

The Operations Overview dashboard provides you with a high-level view of objects which make up your virtual environment. You can view an aggregate of the virtual machine growth trends across the different data centers that vRealize Operations Manager monitors.

You can also view a list of all your data centers with inventory information about how many clusters, hosts, and virtual machines you are running in each of your data centers. By selecting a particular data center, you can narrow down on the areas of availability and performance. The dashboard provides a trend of known issues in each of your data centers based on the alerts which have triggered in the past.

You can also view a list of the top 15 virtual machines in the selected data center which might be contending for resources.

You can use the dashboard widgets in several ways.

- **Environment Summary:** Use this widget to view a summary of the overall inventory of your environment.
- **Select a Datacenter:** Use this widget to select the data center for which you want to view operational information. You can use the filter to narrow your list based on several parameters. After you identify the data center you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Cumulative Up-time of all Clusters:** Use this widget to view the overall health of the clusters in the data center you selected. The metric value is calculated based on the uptime of each ESXi host, when you take into account one host as the HA host. If the number displayed is less than 100%, it means that at least two hosts within the cluster were not operational for that period.
- **Alert Volume (in selected DC):** Use this widget to view the breakdown of alert trends based on their criticality.
- **Top-N:** You can also view a list of 15 VMs that had the highest average CPU contention, the highest use of memory, and the highest disk latency for the last 24 hours. To obtain specific data, you can manually set the time to the time of the problem. To set the time, click the **Edit Widget** icon from the title bar of the widget and edit the **Period Length** drop-down menu.

vSAN Operations Overview

The vSAN Operations Overview dashboard provides an aggregated view of the health and performance of your vSAN clusters.

You can use this dashboard to get a complete view of your vSAN environment and what components make up the environment. You can also view the growth trend of virtual machines served by vSAN.

You can use the dashboard to understand the utilization and performance patterns for each of your vSAN clusters by selecting one from the list that is provided. You can use this dashboard to track vSAN properties such as hybrid or all flash, deduplication and compression, or a stretched vSAN cluster.

You can view the historic performance, utilization, growth trends, and events related to vSAN, with the current state.

Optimize Dashboards

The Optimize group of dashboards include the Optimize Performance, Access Cost, and Optimization History dashboards.

Assess Cost Dashboard

The Assess Cost dashboard gives you cost and reclaimable resources for your data centers and clusters.

The Assess Cost dashboard belongs to the Optimize group of dashboards. This dashboard is ideal for executives, finance, or others who are accountable for overall IT spend. It is also helpful for identifying and planning cost optimization initiatives.

Any cost information shown in this dashboard is using the currency settings you select during vRealize Operations Manager configuration.

The dashboard provides an overview of the cost and inventory for your environment, including total cost of ownership and a total of the potential cost savings based on vRealize Operations capacity engine recommendations.

Individual data centers are listed showing population details, cost information, and reclaimable resources.

At the bottom of the dashboard, you can find the top 10 lists for the most expensive and least expensive clusters in your environment. These lists include the total monthly cost and count of hosts, datastores, and virtual machines. These lists can be helpful in identification of under-utilized clusters by noting the number of virtual machines hosted relative to the monthly cluster cost.

Optimization History Dashboard

The Optimization History dashboard displays the results of optimization activity.

The Optimization History dashboard belongs to the Optimize group of dashboards. The dashboard covers three optimization benefits; optimize performance, optimize capacity, and optimize virtual machine placement.

Optimizing performance can be performed in vRealize Operations Manager using Workload Optimization, or started on demand. The charts on this row show a box for each data center or custom data center and the optimization recommendation. Green indicates an optimized data center or custom data center. A red box means that optimization might be required, and a white box means that optimization is not configured for that object.

For capacity optimization, this row provides a summary of the average VM cost per month, the savings that can be achieved through reclaiming idle or powered off virtual machines, or deleting old snapshots.

Virtual Machine Happiness is a term used to describe VMs that are getting the resources they need, when they need them. You can also see recent vMotion activity related to vSphere's Distributed Resource Scheduler, which together with vRealize Operations predictive DRS feature makes sure your VMs are getting the resources they need. Workload placement vMotions are also shown as Non-DRS Moves in the graph.

Optimize Performance Dashboard

The Optimize Performance dashboard helps you identify virtual machines that can be configured to improve overall performance.

The capacity analytics engine intelligently calculates the settings for CPU and memory for virtual machines to give you the best performance and accurate resource allocation for all workloads.

The dashboard organizes virtual machines by undersized - or virtual machines that are not being served well - and oversized - which are virtual machines that are not using all allocated resources. Both categories consider CPU and memory usage and provide recommendations for optimal sizing.

Performance Troubleshooting Dashboards

The dashboards in the Performance Troubleshooting category cater to the administrators responsible for managing the performance and availability of the virtual machines running in the virtual infrastructure. This category runs you through a guided workflow to answer questions that help you with the troubleshooting process. The dashboards in this category identify and isolate problems that may impact your applications. They provide insight into the full stack to isolate and identify the root cause quickly.

Key questions these dashboards help you answer are as follows:

- Is the application performance impacted due to virtual infrastructure?
- Are noisy neighbors impacting multiple virtual machines and corresponding applications?
- Are there active alerts which require action?
- Are there any known issues impacting the performance and availability of a vSAN cluster?

Troubleshoot a Cluster

The Troubleshoot a Cluster dashboard allows you to identify clusters that have issues and isolate them easily.

You can use the search option to identify a cluster that has an issue. You can also sort the clusters based on the number of active alerts.

After you select the cluster you want to work with, you can view a quick summary of the number of hosts in that cluster and the VMs served by the cluster. The dashboard provides you with current and past utilization trends and also known issues in the cluster in the form of alerts.

You can view the hierarchy of objects related to the cluster and review the status to identify if the objects are impacted because of the current health of the cluster. You can quickly identify any contention issues by looking at the maximum and average contention faced by the VMs on the selected cluster. You can narrow down and view those VMs that have resource contention and take specific steps to troubleshoot and resolve issues.

You can use the dashboard widgets in several ways.

- **Search for a cluster:** Use this widget to select the cluster for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Is your cluster busy?:** Use this widget to view the CPU and memory demand.
- **Are there active alerts on your cluster:** Use this widget to view only the critical alerts.
- **Are the relatives healthy?:** Use this widget to view the hierarchy of the objects related to the cluster and if any of the objects are impacted.
- View the maximum and average CPU, memory, and disk latency for the VMs. If the VM faces contention, it might mean that the underlying infrastructure does not have enough resources to meet the needs of the VMs.
- View a list of VMs that face CPU, memory, and disk latency contention. You can then troubleshoot and take steps to resolve the problem.

Troubleshoot a Datastore

The Troubleshoot a Datastore dashboard allows you to identify storage issues and act on them.

You can use the search option to identify a datastore that has an issue or you can identify a datastore that has high latency as seen in red on the heat map. You can also sort all the datastores with active alerts and troubleshoot the datastore with known issues.

You can select a datastore to see its current capacity and utilization with the number of VMs served by that datastore. The metric charts help you view historical trends of key storage metrics such as latency, outstanding IOs, and throughput.

The dashboard also lists the VMs served by the selected datastore and helps you analyze the utilization and performance trends of those VMs. You can migrate the VMs to other datastores to even out the IO load.

You can use the dashboard widgets in several ways.

- **Search for a datastore:** Use this widget to select the datastore for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the datastore you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Are there active alerts on your datastore:** Use this widget to view only the critical alerts.
- **Are the relatives healthy?:** Use this widget to view the hierarchy of the objects related to the datastore and if any of the objects are impacted.
- **Is your datastore experiencing high latency?** and **Any outstanding disk I/Os?:** Use these widgets to view those datastores with high latency and outstanding disk I/O trends. Ideally, your datastores must not have outstanding disk I/O.
- **How many IOPS is your datastore serving** and **Latency trend for the I/Os done by the VM:** Use these widgets to view the current IOPS and latency of the VMs in the selected datastore.

- Use the other widgets in the dashboard to view trends for the selected datastore regarding disk latency, IOPS, and throughput, VMs served by the datastore and I/O pattern of the selected VM.

Troubleshoot a Host

The Troubleshoot a Host dashboard allows you to search for specific hosts or sort hosts with active alerts. ESXi hosts are the main source of providing resources to a VM and are critical for performance and availability.

To view the key properties of each host, select a host from the dashboard. You can ensure that the host is configured according to the virtual infrastructure design. Any deviation from standards might cause potential issues. You can use the dashboard to answer key questions about current and past utilization and workload trends over the last week. You can also view if the VMs served by the host are healthy.

Since the dashboard lists all the critical events that might affect the availability of the hosts, you can view hardware faults associated with the host. You can view a list of the top 10 VMs that demand CPU and memory resources from the identified host.

Troubleshoot a VM Dashboard

The Troubleshoot a VM dashboard helps an administrator to troubleshoot everyday issues in a virtual infrastructure. While most of the IT issues in an organization are reported at the application layer, you can use the guided workflow in this dashboard to help investigate an ongoing or a suspected issue with the VMs supporting the impacted applications.

You can search for a VM by its name or you can sort the list of VMs with active alerts on them to start your troubleshooting process. When you select a VM, you can view its key properties to ensure that the VM is configured as per your virtual infrastructure design. Any deviation from standards may cause potential issues. You can view known alerts and the workload trend of the VM over the past week. You can also view if any of the resources serving the virtual machine have an ongoing issue.

The next step in the troubleshooting process allows you to eliminate the major symptoms which might impact the performance or availability of a VM. You can use key metrics to find out if the utilization patterns of the VMs are abnormal or if the VM is contending for basic resources such as CPU, memory, or disk.

You can use the dashboard widgets in several ways.

- **Search for a VM:** Use this widget to view all the VMs in the environment. You can select the VM you want to troubleshoot. You can use the filter to narrow your list based on several parameters, such as name, folder name, associated tag, host, or vCenter Server. After you identify the VM you want to troubleshoot, select it. The dashboard is automatically populated with the relevant data.
- **About the VM:** Use this widget to understand the context of the VM. This widget also lends insights to analyze the root cause of the problem or potential mitigations.
- **Are there active alerts on the VM?:** Use this widget to view active alerts. To see noncritical alerts, click the VM object.

- **Is the VM working hard over the last week?:** Use this widget to view the workload trend of the VM for the last week.
- **Are the relatives healthy?:** Use this widget to view the ESXi host where the VM is now running. This host might not be the ESXi host where the VM was running in the past. You can view the remaining related objects and see whether they might contribute to the problem.
- **Is the VMs demand spiking or abnormal?:** Use this widget to identify spikes in the VM demand for any of the resources such as CPU, memory, and network. Spikes in the demand might indicate an abnormal behavior of the VM or that the VM is undersized. The memory utilization is based on the Guest OS metric. It requires VMware Tools 10.0.0 or later and vSphere 6 Update 1 or later. If you do not have these products, the metric remains blank.
- **Is the VM facing contention?:** Use this widget to identify whether the VM is facing contention. If the VM is facing contention, the underlying infrastructure might not have enough resources to meet the needs of the VM.
- **Does the cluster serving the VM have contention?:** Use this widget to view the trend for the maximum CPU contention for a VM within the cluster. The trend might indicate a constant contention within the cluster. If there is contention, you must troubleshoot the cluster as the problem is no longer with the VM.
- **Does the datastore serving the VM have latency?:** Use this widget to help you correlate the latency at the datastore level with the total latency of the VM. If the VM has latency spikes, but the datastore does not have such spikes, it might indicate a problem with the VM. If the datastore faces latency as well, you can troubleshoot to find out why the datastore has these spikes.
- **Parent Host and Parent Cluster:** Use these widgets to view the host and the cluster on which the VM resides.

Troubleshoot vSAN Dashboard

The Troubleshoot vSAN dashboard helps you view the properties of your vSAN cluster and the active alerts on the cluster components. The cluster components include hosts, disk groups, or the vSAN datastores.

You can select a cluster from the dashboard and then list all the known problems with the objects associated with the cluster. The objects include clusters, datastores, disk groups, physical disks, and VMs served by the selected vSAN cluster.

You can view the key use and performance metrics from the dashboard. You can also view the usage and performance trend of the cluster for the last 24 hours. You can also view historical issues and analyze the host, disk group, or physical disk.

You can use the heat maps within the dashboard to answer questions about write buffer usage, cache hit ratio, and host configurations. You can also use the heat maps to answer questions about physical issues with capacity and cache disks, such as drive wear out, drive temperature, and read-write errors.

You can use the dashboard widgets in several ways.

- **Search for a vSAN cluster:** Use this widget to search vSAN clusters. You can view the details of each vSAN cluster including the number of hosts, VMs, cache disks, capacity disks, and cluster type are provided. You can also view if the vSAN cluster is dedupe and compression enabled, and stretched.
- **Any alerts on the cluster, hosts, VMs or disks?:** Use this widget to view alerts on the cluster, VMs, or disks in your environment.
- **Are the relatives healthy?:** Use this widget to view the health, risk, and efficiency of the relatives. This widget also allows you to view the health of the datastore in a host and disks in each disk group.
- **Are outstanding I/Os high?:** Use this widget to view the key performance metrics. The widget indicates outstanding I/Os within 24 hours time period.
- **Are VMs facing read latency?:** Use this widget to view the read latency of VMs.
- **Are VMs facing write latency?:** Use this widget to view the write latency of VMs.
- **Is the write buffer low?:** Use this widget to view the usage of the write buffer on diskgroups in a cluster.
- **Are the hosts consistently configured?:** Use this widget to view the participating hosts in the selected cluster and to determine if the hosts are consistently configured.
- **Cache Disks: Any hardware issues?:** Use this widget to view the individual cache disks measured against various metrics.
- **Capacity Disks: Any hardware issues?:** Use this widget to view the individual capacity disks measured against various metrics.

Troubleshoot with Logs Dashboard

When vRealize Operations Manager is integrated with vRealize Log Insight, you can access the custom dashboards and content pack dashboards from the Troubleshoot with Logs dashboard. You can view graphs of log events in your environment, or create custom sets of widgets to access the information that matters most to you.

You can investigate an ongoing issue within your virtual infrastructure using the logs. You can view predefined views created within vRealize Log Insight to answer questions from predefined queries within vRealize Log Insight.

You can correlate metrics and queries within vRealize Operations Manager to troubleshoot issues across applications and infrastructure.

For more information about the Troubleshoot with Logs dashboard, see the [vRealize Log Insight documentation](#).

To access the Troubleshoot with Logs dashboard from vRealize Operations Manager, you must either:

- Configure the vRealize Log Insight adapter from the vRealize Operations Manager interface, or

- Configure vRealize Operations Manager in vRealize Log Insight.

For more information on configuring, see [Configuring vRealize Log Insight with vRealize Operations Manager](#).

vRealize Automation Dashboards

With the vRealize Automation dashboards, you can monitor and troubleshoot objects in your cloud infrastructure.

The following vRealize Automation solution dashboards are added to the predefined vRealize Operations Manager dashboards:

- Application Overview
- Environment Overview
- Resource Consumption Overview
- Top-N

Application Overview Dashboard

You can use the widgets in the Application Overview dashboard to view the blueprint objects and the blueprint deployment details.

You can use the Application Overview dashboard to view the hierarchy, the properties of the blueprint and deployments, and the metric information.

You can use the dashboard widgets in several ways.

- **Blueprint List:** Use this widget to view the blueprint objects in the environment.
- **Blueprint Overview:** Use this widget to view the relationship between the blueprint objects and the deployment, virtual machines, cluster compute resources, and the datastore objects. To find the deployment, virtual machine, and other related details, click the blueprint object.
- **Blueprint Property List:** Use this widget to view the properties of the blueprint object such as the total cost, average deployment time, and the average cost of the blueprint object .
- **Deployment List:** Use this widget to view the blueprint objects deployed in the environment.
- **Deployment Property List:** Use this widget to view the properties for the deployment object such as the cost until date and the approval time for each deployment.
- **Blueprint Deployment Info:** Use this widget to select a metric. You can view the details in the Metric Chart widget.
- **Metric Chart:** Use this widget to view the relevant data based on the metric you select in the Blueprint Deployment Info widget.

Environment Overview Dashboard

You can use the Environment Overview dashboard to view information about the tenants and the related alerts.

You can use the Environment Overview dashboard to perform some of the following tasks:

- To view the active alerts on vCenter resources that are managed by vRealize Automation.

You can use the dashboard widgets in several ways.

- **Environment Summary.** Use this widget to view the health of tenants, business groups, virtual machines, blueprints, reservations, deployments, cluster compute resources and the relationships between these objects. If you double-click an object in the Environment Overview widget, you can view detailed information for the object.
- **Tenant List.** Use this widget to view the tenant objects available in the environment. You can see a data grid with a list of objects in the inventory on which you can sort and search.
- **Business Group List.** Use this widget to view the business group objects available in the environment. You can see a data grid with a list of objects in the inventory on which you can sort and search. You can see a data grid with a list of objects in the inventory on which you can sort and search.
- **Configured Users.** Use this widget to view the business group name and the user configured for the business group.
- **vRealize Automation Inventory.** Use this widget to view the objects available for each vRealize Automation solution that is deployed in the environment.
- **vRealize Automation Managed Clusters.** Use this widget to view the vCenter clusters which are managed by vRealize Automation. You can see a data grid with a list of objects in the inventory on which you can sort and search.
- **Top Alerts.** Alerts with the greatest significance on the selected objects it is configured to monitor. The top alerts include a short description of alerts configured for the widget. The alert name opens a secondary window from which you can link to the alert details. In the alert details, you can begin resolving the alerts.

Resource Consumption Overview Dashboard

You can use the widgets in the Resource Consumption Overview dashboard to view the resources consumed by vRealize Automation on a vCenter Server.

You can use the Resource Consumption Overview dashboard widgets in several ways.

- **Tenant List:** Use this widget to view the tenant objects available in the environment. You can see a data grid with a list of tenants objects in the inventory on which you can sort and search.
- **Business Group List:** Use this widget to view the business group objects available in the environment. You can see a data grid with a list of objects in the inventory on which you can sort and search.
- **Reservation List:** Use this widget to view the reservation objects available in the environment. You can see a data grid with a list of objects in the inventory on which you can sort and search.
- **Tenant Capacity:** Use this widget to analyze the capacity of the tenant object.
- **Business Group Capacity:** Use this widget to view the memory, storage, and quota capacity that is allocated, reserved, and free for each business group object.

- **Reservation Capacity:** Use this widget to view the memory, storage, and quota capacity that is allocated, reserved, and free for each reservation object.
- **Tenant Memory Trend:** Use this widget to view and analyze a seven-day trend for the memory allocated, reserved, and free for a tenant object.
- **Tenant Storage Trend:** Use this widget to view and analyze a seven-day trend for the storage allocated, reserved, and free for a tenant object.

Top-N Dashboard

You can use the widgets in the Top-N dashboard to view the top results from analysis of blueprints, business groups, and tenants that you select.

You can use the Top-N dashboard to perform some of the following tasks:

- To view the most popular blueprints, business groups, and tenants.
- To view the business groups that have the most critical alerts.

You can use the dashboard widgets in several ways.

- **Tenant with most critical alerts.** Use this widget to view the top- five tenant objects that have the most critical alerts.
- **Business Groups with most Critical Alerts.** Use this widget to view the top-five business group objects that have the most critical alerts.
- **Tenant with most failed requests.** Use this widget to view the top-five tenant objects that have the most failed requests.
- **Most popular deployed Tenant.** Use this widget to view the top-five most popular deployed tenant objects in the environment.
- **Most popular deployed Business Group.** Use this widget to view the top-five most popular deployed business group objects in the environment.
- **Most Popular Deployed Blueprints.** Use this widget to view the top-five most popular deployed blueprint objects in the environment.
- **Most Popular Deployed Business Group (7 day trend).** Use this widget to view graphical trends that contain metrics for the virtual machine count that has been deployed the most for the business group object over a seven-day period.
- **Most Popular Deployed Blueprints (7 day trend).** Use this widget to view graphical trends that contain metrics for the virtual machine count that has been deployed the most for the blueprint object over a seven-day period.

Create and Configure Dashboards

To view the status of all objects in vRealize Operations Manager, create a dashboard by adding widgets or views. You can create and modify dashboards and configure them to meet your environment needs.

Procedure

- 1 In the menu, click **Dashboards**.
- 2 Click **Actions > Create Dashboard** to create and configure a dashboard.
- 3 Complete the following steps to:
 - a Enter a name for the dashboard.
[Dashboard Name](#)
 - b Add widgets or views to the dashboard.
[Widget or View List Details](#)
 - c Configure widget interactions.
[Widget and View Interactions Details](#)
 - d Create dashboard navigation.
[Dashboard Navigation Details](#)
- 4 Click **Save**.
- 5 Click **Actions > Edit Dashboard** to modify the dashboard.

Dashboard Name

The name and visualization of the dashboard as it appears on the vRealize Operations Manager Home page.

Where You Add a Name in a Dashboard

To create or edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. Enter a name in the **New Dashboard** field.

If you use a forward slash while entering a name, the forward slash acts as a group divider and creates a folder with the specified name in the dashboards list if the name does not exist. For example, if you name a dashboard **cclusters/hosts**, the dashboard is named hosts under the group cclusters.

Widget or View List Details

vRealize Operations Manager provides a list of widgets or views that you can add to your dashboard to monitor specific metrics and properties of objects in your environment.

Where You Add Widgets or Views to a Dashboard

To create or edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard.

How to Add Widgets or Views to a Dashboard

In the widgets list panel, you see a list of all the predefined vRealize Operations Manager widgets or views. Drag the widget or view to the dashboard workspace in the upper panel.

To locate a widget or view, you can type the name or part of the name of a widget or view in the **Filter** option. For example, when you enter **top**, the list is filtered to display the Top Alerts, Top-N, and Topology Graph widgets. You can then select the widget you require.

Most widgets or views must be configured individually to display information. For more information about how to configure each widget, see [Widgets](#).

How to Arrange Widgets or Views in a Dashboard

You can modify your dashboard layout to suit your needs. By default, the first widgets or views that you add are automatically arranged horizontally wherever you place them.

- To position a widget or a view, drag the widget or view to the desired location in the layout. Other widgets and views automatically rearrange to make room.
- To resize a widget or a view, drag the bottom right corner of the widget or the view.

Widget and View Interactions Details

You can connect widgets and views so that the information they show depends on each other.

Where You Create Widget and View Interactions

To create interactions for widgets or views in a dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. From the toolbar, click **Show Interactions**.

How to Create and Remove Widget Interactions

The list of available interactions depends on the widgets or views in the dashboard. Widgets and views can provide, receive, and can both provide and receive interactions at the same time.

To create interactions, click **Show Interactions**. Click a provider plug and drag to the receiver. You can also apply interactions from receiver to provider plugs. For more information about how interactions work, see [Widget Interactions](#).

To remove interactions, click on the interaction line and select **Remove Interaction**. You can also click the provider plug and select **Remove Interaction > <widget name>**.

Dashboard Navigation Details

You can apply sections or context from one dashboard to another. You can connect widgets and views to widgets and views on other dashboards to investigate problems or better analyze the provided information.

Where You Add Another Dashboard

To create dashboard navigation to a dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. In the dashboard workspace, click **Show Interactions**. From the **Select Another Dashboard** drop-down menu, select the dashboard to which you want to navigate.


How Dashboard Navigation Works

You can create dashboard navigation only for provider widgets and views. The provider widget or view sends information to the destination widget or view. When you create dashboard navigation, the destination widgets or views are filtered based on the information type they can receive.

How to Add a Dashboard Navigation to a Dashboard

The list of available dashboards for navigation depends on the available dashboards and the widgets and views in the current dashboard. To add navigation, you can drag and drop from a sender widget interaction plug to a receiver widget interaction plug. You can select more than one applicable widget or view.

Note If a dashboard is unavailable for selection, it is unavailable for dashboard navigation.

The Dashboard Navigation icon () appears in the top menu of each widget or view when a dashboard navigation is available.

Managing Dashboards

You can change the order of the dashboard tabs, configure vRealize Operations Manager to switch from one dashboard to another, create dashboard folders to group the dashboards in a way that is meaningful to you, share a dashboard or dashboard template with one or more user groups, and transfer selected dashboards to a new owner.

Reorder and Switch Dashboards

You can change the order of the dashboard tabs on your home page. You can configure vRealize Operations Manager to switch from one dashboard to another. This feature is useful if you have several dashboards that show different aspects of your enterprise's performance and you want to look at each dashboard in turn.

Where You Configure a Dashboard Order and Automatic Switch

To reorder and configure a dashboard switch, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Reorder/Autoswitch Dashboards**.

How You Reorder the Dashboards

The list shows the dashboards as they are ordered. Drag the dashboards up and down to change their order on the home page.

How You Configure an Automatic Dashboard Switch

- 1 Double-click a dashboard from the list to configure.
- 2 From the Auto Transition drop-down menus, select **On**.
- 3 Select the switch time interval in seconds.
- 4 Select the dashboard to switch to and click **Update**.
- 5 Click **Save** to save your changes.

On the home page, the current dashboard will switch to the dashboard that is defined after the specified time interval.

Manage Summary Dashboards

The **Summary** tab provides you with an overview of the state of the selected object, group, or application. You can change the **Summary** tab with a dashboard to get information specific to your needs.

Where You Configure a Summary Tab Dashboard

To manage the summary dashboards, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Manage Summary Dashboards**.

How You Manage the Summary Tab Dashboard

Table 7-5. Manage Summary Dashboards Options

Option	Description
Adapter Type	Adapter type for which you configure a summary dashboard.
Filter	Use a word search to limit the number of adapter types that appear in the list.
Name	List with all available objects.
Use Default icon	Click to use vRealize Operations Manager default Summary tab.
Detail Page	Shows what kind of Summary tab you use for the selected object.
Assign a Dashboard icon	Click to view the Dashboard List dialog box that lists all the available dashboards.

To change the Summary tab for an object, select the object in the left panel, click the **Assign a Dashboard** icon. Select a dashboard for it from the Dashboard List dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard you have associated to the object type when you navigate to the **Summary** tab of the object details page.

Manage Dashboard Groups

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

Where You Configure a Dashboard Group

To manage the dashboard groups, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Manage Dashboard Groups**.

How You Manage the Dashboard Groups

Table 7-6. Manage Dashboard Groups Options

Option	Description
Dashboard Groups	A hierarchy tree with all available group folders.
Dashboards List	A list with all available dashboards.

To create a dashboard group folder, right-click the **Dashboard Groups** folder or another folder and click **Add**. To add a dashboard, drag one from the Dashboards list to the folder.

Share Dashboards with Users

You can share a dashboard or dashboard template with one or more user groups. When you share a dashboard, it becomes available to all the users in the user group that you select. The dashboard appears the same to all the users who share it. If you edit a shared dashboard, the dashboard changes for all users. Other users can only view a shared dashboard. They cannot change it.

Where You Share a Dashboard From

To share a dashboard, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Share Dashboards**.

Table 7-7. Share Dashboards Options

Option	Description
Accounts Group	All available groups with which you can share a dashboard.
Shared Dashboards	All available dashboards and templates that you can share. You can switch between dashboard tabs and dashboard templates by clicking the Share Dashboard Tabs/Templates icon.

How You Manage a Shared Dashboard Tab

To share a dashboard tab, navigate to the dashboard in the list of Shared Dashboards and drag it to the group to share it with on the left.

To stop sharing a dashboard with a group, click that group on the left panel, navigate to the dashboard in the right panel, and click the **Stop Sharing** icon above the list.

To stop sharing a dashboard with more than one group, click the **Not Grouped** name on the left panel, navigate to the dashboard in the right panel, and click the **Stop Sharing** icon above the list.

Options for Sharing Dashboards

You can share predefined or custom dashboards using URLs, emails, and by copying the code to embed the dashboard into confluence or other internal official web pages. You can also assign and unassign a dashboard to specific user groups and export the dashboard configuration details.

When you use a non-authenticated shared URL, as a user you can open the dashboard in a new browser session. If you have already logged into vRealize Operations Manager in another session, you are redirected to this dashboard and the user authentication permissions apply. To ensure that the non-authenticated URL opens the intended dashboard, as a user you must log out from all existing user sessions.

The dashboard shared with the URL opens in a page where you can access all the widgets within the dashboard and you can interact with the given widgets at the same time. A non-authenticated dashboard however, does not allow you to browse to other areas of vRealize Operations Manager.

Where You Can Access the Options to Share Dashboards

From the menu, select **Dashboards**. Click on an existing dashboard and then click the **Share Dashboard** icon in the top right corner.

Table 7-8. Options in the Share Dashboard Dialog Box

Option	Description
URL	<p>Allows you to copy the tiny URL for the selected dashboard.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 Months, or Never Expire. ■ Click Copy Link to copy the link to a new window from where you can view the dashboard. <hr/> <p>Note</p> <ul style="list-style-type: none"> ■ As a user, if you open a shared link and you are logged into vRealize Operations Manager, you are navigated to your default dashboard, instead of viewing the shared one. ■ As a user, if you log in to the same IP that was shared with you previously, you cannot access the page with the same browser. ■ As a user, ensure that you have the following permission: Dashboards > Dashboard Management > Share (Public). <hr/> <p>You can stop sharing a dashboard you had previously shared. To stop sharing a dashboard, click the Unshare Link option and enter the URL of the dashboard that you want to stop sharing and click Unshare.</p> <p>Authentication is not required to view the shared dashboard.</p>
Email	<p>Allows you to send an email with the URL details of the dashboard, to a specific person.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 months, or Never Expire. ■ Configure an SMTP instance. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts. ■ Enter an email address and click the Send Email button to send an email with the URL details of the dashboard. <p>Authentication is not required to view the shared dashboard.</p>

Table 7-8. Options in the Share Dashboard Dialog Box (continued)

Option	Description
Embed	<p>Provides an embedded code for the dashboard. You can use this code to embed the dashboard in relevant confluence pages that your company executives routinely use and analyze.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 Months, or Never Expire. <hr/> <p>Note</p> <ul style="list-style-type: none"> ■ If you embed a dashboard in the Text widget, the widget does not display any data. ■ When you open an HTML/confluence page with an embedded dashboard from the same browser that you have logged into vRealize Operations Manager, the dashboard does not load. <hr/> <p>Authentication is not required to view the shared dashboard.</p>
Groups	<p>Allows you to assign and unassign a dashboard to specific user groups.</p> <ul style="list-style-type: none"> ■ Select the group to which you want to grant dashboard access from the drop-down menu and click Include. You can include more than one dashboard. ■ From the label, select the cross mark to unassign the dashboard. <p>Log in to vRealize Operations Manager to view the shared dashboard.</p>
Export	<p>Allows you to export the dashboard configuration details. Log in to vRealize Operations Manager to export/import a dashboard.</p>

Views

vRealize Operations Manager provides several types of views. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective. vRealize Operations Manager Views also show information that the adapters in your environment provide.

You can configure vRealize Operations Manager views to show transformation, forecast, and trend calculations.

- The transformation type determines how the values are aggregated.
- The trend option shows how the values tend to change, based on the historical, raw data. The trend calculations depend on the transformation type and roll up interval.
- The forecast option shows what the future values can be, based on the trend calculations of the historical data.



Create Views

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_create_view_vrop)

You can use vRealize Operations Manager views in different areas of vRealize Operations Manager.

- To manage all views, in the menu, click **Dashboards**, and then in the left pane click **Views**.
- To see the data that a view provides for a specific object, navigate to that object, click the **Details** tab, and click **Views**.
- To see the data that a view provides in your dashboard, add the View widget to the dashboard.
- To have a link to a view in the Further Analysis section, select the Further Analysis option on the view workspace visibility step.

Views and Reports Ownership

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Views Overview

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective.

In the menu, click **Dashboards**, and then in the left pane click **Views** to access the Views page.

On the Views page you can create, edit, delete, clone, export, and import views.

You can order the listed views by name, type, description, subject, or owner.

You can limit the views list by adding a filter from the upper-right corner of the panel.

Table 7-9. Filter Groups

Filter Group	Description
Name	Filter by the view name. For example, type my view to list all views that contain the my view phrase in their name.
Type	Filter by the view type.
Description	Filter by the view description. For example, type my view to list all views that contain the my view phrase in their description.
Subject	Filter by the subject.
Owner	Filter by the owner.

Views and Reports Ownership

The owner of views, reports, or templates might change over time.

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Create and Configure a View

To collect and display information for a specific object, you can create a custom view.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Click the **Create View** icon to create a view.
- 3 Complete the steps in the left pane to:
 - a Enter a name and description for the view.
[Name and Description Details](#)
 - b Change the presentation of a view.
[Presentation Details](#)

- c Select the base object type for a view.

[Subjects Details](#)

- d Add data to a view.

[Data Details](#)

- e Change the visibility of a view.

[Visibility Details](#)

- 4 Click **Save**.

- 5 From the Views page, click the **Edit View** icon to modify the view.

Name and Description Details

The name and description of the view as they appear in the list of views on the Views page.

To add a name and description to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Name and Description**.

Table 7-10. Name and Description Options in the View Workspace

Option	Description
Name	Name of the view as it appears on the Views page.
Description	Description of the view.

Presentation Details

A presentation is a way the collected information for the object is presented. Each type of view helps you to interpret metrics and properties from a different perspective.

To change the presentation of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Presentation**. If you create a view, complete the required previous steps.

Table 7-11. Presentation Options in the View Workspace

View Type	Description
List	Provides tabular data about specific objects in the monitored environment. Column count is limited to 25 in a PDF report and 50 in a CSV report. Page count is unlimited.
Summary	Provides tabular data about the use of resources in the monitored environment.
Trend	Uses historic data to generate trends and forecasts for resource use and availability in the monitored environment.

Table 7-11. Presentation Options in the View Workspace (continued)

View Type	Description
Distribution	<p>Provides aggregated data about resource distribution in the monitored environment.</p> <p>When you add a distribution type of View to a dashboard, you can click a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment.</p>
Text	<p>Inserts the provided text. The text can be dynamic and contain metrics and properties.</p> <p>You can format text to increase or decrease the font size, change the font color, highlight text, and align text to the left, right, or center. You can also make the selected text appear bold, in italics, or underlined.</p> <p>By default the text view is available only for report template creation and modification. You can change this on the Visibility step of the view workspace.</p>
Image	<p>Inserts a static image.</p> <p>By default the image view is available only for report template creation and modification. You can change this on the Visibility step of the view workspace.</p>

You can see a live preview of the view type when you select a subject and data, and **Select preview source**.

How to Configure the Presentation of a View

Some of the view presentations have specific configuration settings.

Table 7-12. Presentation Configuration Options in the View Workspace

View Type	Configuration Description
List	<ul style="list-style-type: none"> ■ Select the number of items per page. Each item is one row and its metrics and properties are the columns. ■ Select the top results. Restricts the number of results. For example, if you list all the clusters in a View, selecting 10 in this option displays the top 10 clusters with the relevant information. You can reduce the number of rows for the purposes of reporting.
Summary	Select the number of items per page. Each row is an aggregated metric or property.

Table 7-12. Presentation Configuration Options in the View Workspace (continued)

View Type	Configuration Description
Trend	<p>Enter the maximum number of plot lines. Limits the output in terms of the objects displayed in the live preview of the view type on the left upper pane. The number you set as the maximum number of plot lines determines the plot lines.</p> <p>For example, if you plot historical data and set the maximum at 30 plot lines, then 30 objects are displayed. If you plot historical, trend, and forecast lines, and set the maximum to 30 plot lines, then only 10 objects are displayed as each object has three plot lines.</p>
Distribution	<p>Select the visualization of the distribution information in a pie chart or a bar chart.</p> <p>Select the distribution type, and configure the buckets count and size.</p> <p>To understand vRealize Operations Manager distribution type, see View Distribution Type.</p>

Coloring

Configuration Option	Description
Colorize	The colors of the slices in the pie chart are displayed in the order of the colors in the color palette.
Select Color	Select the color that you want the chart to appear in. If there is more than one slice in a pie chart, the colors are chosen sequentially from the color palette. In a bar chart, the bars are all the same color.

Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Dynamic distribution

You specify in details how vRealize Operations Manager distributes the data in the buckets.

Table 7-13. Dynamic Distribution Configuration Options

Configuration Option	Description
Buckets Count	The number of buckets to use in the data distribution.
Buckets Size Interval	The bucket size is determined by the defined interval divided by the specified number of buckets.

Table 7-13. Dynamic Distribution Configuration Options (continued)

Configuration Option	Description
Buckets Size Logarithmic bucketing	The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data.
Buckets Size Simple Max/Min bucketing	The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets.

Manual distribution

You specify the number of buckets and the minimum and maximum values of each bucket.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distribute the data.

View Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Visualization

You can choose to view the data as a pie chart or a bar chart. When you add a distribution type of View to a dashboard, you can click on a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment. You can select the display colors for single or multi-colored charts.

Dynamic distribution

You specify in details how vRealize Operations Manager distributes the data in the buckets.

Table 7-14. Dynamic Distribution Configuration Options

Configuration Option	Description
Buckets Count	The number of buckets to use in the data distribution.
Buckets Size Interval	The bucket size is determined by the defined interval divided by the specified number of buckets.
Buckets Size Logarithmic bucketing	The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data.
Buckets Size Simple Max/Min bucketing	The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets.

Manual distribution

You specify the number of buckets and the minimum and maximum values of each bucket.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distribute the data.

If you increase the number of buckets, you can see more detailed data.

Subjects Details

The subject is the base object type for which the view shows information.

To specify a subject for a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Subjects**. If you create a new view, complete the required previous steps.

The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. You can limit the level where the view appears with the Blacklist option in the **Visibility** step.

View availability depends on the view configuration subject, inventory view, user permissions, and view Visibility settings.

For list views with **Symptom** as a subject, the following columns can be sorted: Criticality Level, Status, Object Type, Object Name, Created on, and Canceled on. You cannot sort the Triggered On and Violation Info columns. If other symptom metrics exist, you cannot sort any of the columns.

Views Applicability

List View

When you navigate through the environment tree, you can see the List view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the List view might be missing at the object containers. For example, you create a List view with subject Host System. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Hosts and Clusters > vSphere World**. Select a vCenter Server, and click the **Details** tab, you can see your List view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Storage > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your List view is missing. Your List view with subject Host System is missing because the object Host System is not included in the vSphere Storage inventory view.

Summary View

When you navigate through the environment tree, you can see the Summary view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the Summary view might be missing at the object containers. For example, you create a Summary view with subject Datastore. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Storage >**

vSphere World. Select a vCenter Server, and click the **Details** tab, you can see your List view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Networking > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your Summary view is missing. Your Summary view with subject datastore is missing because the object Datastore is not included in the vSphere Networking inventory view.

Trend View

When you navigate through the environment tree, you can see the Trend view only at the subjects that you specify during the view configuration. For example, you create a Trend view with subject Virtual Machine. When you navigate to a virtual machine in the navigation tree, you see your view.

Distribution View

When you navigate through the environment tree, you can see the Distribution view only at the object containers of the subjects that you specify during the view configuration. Depending on the inventory view, the Distribution view might be missing at the object containers. For example, you create a Distribution view with subject Host System. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Hosts and Clusters > vSphere World**. Select a vCenter Server, and click the **Details** tab, you can see your Distribution view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Networking > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your Distribution view is missing. Your Distribution view with subject Host System is missing because the object Host System is not included in the vSphere Networking inventory view.

Text View

When you navigate through the environment tree, you can see the Text view only at the subjects that you specify during the view configuration. For example, you create a Text view with subject vCenter Server. When you navigate to a vCenter Server in the navigation tree, you see your view. If you did not specify a subject, you see your view for every subject in the environment.

Image View

The Image view is applicable for every object in the environment.

Note Views applicability depends also on your user permissions and the view Visibility configuration.

Views Applicability

Views might not always appear where you expect them to. The main applicability of views depends on the view subject and the inventory view.

List View

When you navigate through the environment tree, you can see the List view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the List view might be missing at the object containers. For example, you create a List view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Storage > vSphere World**, select the same vCenter Server, and click the **Details** tab, your List view is missing. Your List view with subject Host System is missing because the object Host System is not included in the vSphere Storage inventory view.

Summary View

When you navigate through the environment tree, you can see the Summary view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the Summary view might be missing at the object containers. For example, you create a Summary view with subject Datastore. When you go to **Environment > vSphere Storage > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Summary view is missing. Your Summary view with subject datastore is missing because the object Datastore is not included in the vSphere Networking inventory view.

Trend View

When you navigate through the environment tree, you can see the Trend view only at the subjects that you specify during the view configuration. For example, you create a Trend view with subject Virtual Machine. When you navigate to a virtual machine in the navigation tree, you see your view.

Distribution View

When you navigate through the environment tree, you can see the Distribution view only at the object containers of the subjects that you specify during the view configuration. Depending on the inventory view, the Distribution view might be missing at the object containers. For example, you create a Distribution view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your Distribution view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Distribution view is missing. Your Distribution view with subject Host System is missing because the object Host System is not included in the vSphere Networking inventory view.

Text View

When you navigate through the environment tree, you can see the Text view only at the subjects that you specify during the view configuration. For example, you create a Text view with subject vCenter Server. When you navigate to a vCenter Server in the navigation tree,

you see your view. If you did not specify a subject, you see your view for every subject in the environment.

Image View

The Image view is applicable for every object in the environment.

Note Views applicability depends also on your user permissions and the view Visibility configuration.

Data Details

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which vRealize Operations Manager collects, calculates, and presents the information for the view.

To add data to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data**. If you create a new view, complete the required previous steps.

How to Add Data to a View

If you selected more than one subject, specify the subject for which you add data. Double-click the data from the tree in the left panel to add it to the view. For each subject the data available to add might be different.

How to Configure the Data Transformation

The data configuration options depend on the view and data type that you select. Most of the options are available for all views.

Table 7-15. Data Configuration Options

Configuration Option	Description
Metric name	Default metric name. Available for all views.
Metric label	Customizable label as it appears in the view or report. Available for all views.
Units	Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU Demand(MHz) from the Units drop-down menu, you can change the value to Hz, KHz, or GHz. If you select Auto , the scaling is set to a meaningful unit. Available for all views.
Sort order	Orders the values in ascending or descending order. Available for List view and Summary view.

Table 7-15. Data Configuration Options (continued)

Configuration Option	Description
Transformation	<p>Determines what calculation method is applied on the raw data. You can select the type of transformation:</p> <ul style="list-style-type: none"> ■ Minimum. The minimum value of the metric over the selected time range. ■ Maximum. The maximum value of the metric over the selected time range. ■ Average. The mean of all the metric values over the selected time range. ■ Sum. The sum of the metric values over the selected time range. ■ First. The first metric value for the selected time range. ■ Last. The last value of a metric within the selected time range. If you have selected Last as the transformation in versions before vRealize Operations Manager 6.7, and the end of specified time range is not before the last five minutes, use the Current transformation. ■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null. ■ Standard Deviation. The standard deviation of the metric values. ■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum. ■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range. ■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on. ■ Expression. Allows you to construct a mathematical expression over already existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, sum/((max + min)/2). You can use the operands of some of the existing transformations such as, max, min, avg, sum, first, last, current. You cannot use standard deviation, forecast, metric correlation, and percentile. <p>Available for all views, except Trend.</p>
Ranges for metric coloring	<p>You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the Red Bound field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats.</p>

Table 7-15. Data Configuration Options (continued)

Configuration Option	Description
Data Series	<p>You can select whether to include historical data, trend of historical data, and forecast for future time in the trend view calculations.</p> <p>Available for Trend view.</p>
Series Roll up	<p>The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.</p> <p>This option is applicable to the Transformation configuration option.</p> <p>Available for all views.</p>
Threshold Lines	<p>You can set a threshold for a single metric:</p> <ul style="list-style-type: none"> ■ None. You have not set a threshold. ■ By Symptom Definition. You can set a threshold value based on a symptom definition. ■ Custom. You can set the threshold value as Warning, Critical, or Immediate. These options are available only for the Custom option. <p>Available for Trend view.</p>

How to Configure Time Settings

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data.

Table 7-16. Time Settings Options

Configuration Option	Description
Time Range Mode	<p>In Basic mode you can select date ranges.</p> <p>In Advanced mode you can select any combination of relative or specific start and end dates.</p>
Relative Date Range	<p>Select a relative date range of data transformation.</p> <p>Available in Basic mode.</p>
Specific Date Range	<p>Select a specific date range of data transformation.</p> <p>Available in Basic mode.</p>

Table 7-16. Time Settings Options (continued)

Configuration Option	Description
Absolute Date Range	<p>Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.</p> <p>The units of time available are: Hours, Days, Weeks, Months, and Years.</p> <p>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.</p> <p>Available in Basic mode.</p>
Relative Start Date	<p>Select a relative start date of data transformation.</p> <p>Available in Advanced mode.</p>
Relative End Date	<p>Select a relative end date of data transformation.</p> <p>Available in Advanced mode.</p>
Specific Start Date	<p>Select a specific start date of data transformation.</p> <p>Available in Advanced mode.</p>
Specific End Date	<p>Select a specific end date of data transformation.</p> <p>Available in Advanced mode.</p>
Currently selected date range	<p>Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM.</p>

How to Break Down Data

You can break down data in List views by adding interval or instance breakdown columns from the **Group By** tab.

Table 7-17. Group By Options

Option	Description
Add interval breakdown column (see data for column settings)	<p>Select this option to see the data for the selected resources broken down in time intervals.</p> <p>In the Data tab, select Interval Breakdown to configure the column. You can enter a label and select a breakdown interval for the time range.</p>
Add instance breakdown column (see data for column settings)	<p>Select this option to see the data for all instances of the selected resources.</p> <p>In the Data tab, select Instance Name to configure the column. You can enter a label and select a metric group to break down all the instances in that group. Deselect Show non-instance aggregate metric to display only the separate instances. Deselect Show only instance name to display the metric group name and instance name in the instance breakdown column.</p> <p>For example, you can create a view to display CPU usage by selecting the metric CPU:0 Usage. If you add an instance breakdown column, the column CPU:0 Usage displays the usage of all CPU instances on separate rows (0, 1, and so on). To avoid ambiguity, you can change the metric label of CPU:0 Usage to Usage.</p>

How to Add a Filter

The filter option allows you to add additional criteria when the view displays too much information. For example, a list view shows information about the health of virtual machines. From the **Filter** tab you add a risk metric less than 50%. Then the view shows the health of all virtual machines with risk less than 50%.

To add filter to a view, select **Content > Views** in the left pane. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data** and click the **Filter** tab in the main panel. If you create a view, complete the required previous steps.

Each subject has a separate filter box. For Alerts Rollup, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 7-18. Filter Add Options

Option	Description
Add	Adds another criteria to the criteria set. The filter returns results that match all of the specified criteria.
Add another criteria	Adds another criteria set. The filter returns results that match one criteria set or another.

How to Add a Summary Row or Column to a View

The summary option is available only for List and Summary views. It is mandatory for the Summary views. You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, select **Content > Views** in the left pane. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data** and click the **Summary** tab in the main panel. If you create a new view, complete the required previous steps.

For the List view, the summary row shows aggregated information by the specified subjects.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

Visibility Details

The view visibility defines where you can see a view in vRealize Operations Manager.

To change the visibility of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Visibility**. If you create a new view, complete the required previous steps.

Table 7-19. View Workspace Visibility Options

Option	Description
Availability	Select where in vRealize Operations Manager you want to see this view. If you want to have the view available in a dashboard, select the check box, add the View widget, and configure it. You can also make the view available in report templates and in the Detail tab of a specific object when you select the specific check box.
Further Analysis	Select the Compliance check box to make the view available in the Compliance tab for a specific object.
Blacklist	Select a subject level where you do not want to see this view. For example, you have a list view with subject virtual machines. It is visible when you select any of its parent objects. You add datacenter in the banned list. The view is not visible anymore on datacenter level.

Editing, Cloning, and Deleting a View

You can edit, clone, and delete a view. Before you do, familiarize yourself with the consequences of these actions.

When you edit a view, all changes are applied to the report templates that contain it.

When you clone a view, the changes that you make to the clone do not affect the source view.

When you delete a view, it is removed from all the report templates that contain it.

User Scenario: Create, Run, Export, and Import a vRealize Operations Manager View for Tracking Virtual Machines

As a virtual infrastructure administrator, you use vRealize Operations Manager to monitor several environments. You must know the number of virtual machines on each vCenter Server instance. You define a view to gather the information in a specific order and use it on all vRealize Operations Manager environments.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

You will create a distribution view and run it on the main vRealize Operations Manager environment. You will export the view and import it in another vRealize Operations Manager instance.

Procedure

1 [Create a vRealize Operations Manager View for Supervising Virtual Machines](#)

To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

2 [Run a vRealize Operations Manager View](#)

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

3 [Export a vRealize Operations Manager View](#)

To use a view in another vRealize Operations Manager, you export a content definition XML file.

4 [Import a vRealize Operations Manager View](#)

To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Create a vRealize Operations Manager View for Supervising Virtual Machines

To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Click the plus sign to create a new view.
- 3 Enter **Virtual Machines Distribution**, the name for the view.

- 4 Enter a meaningful description for the view.

For example, **A view showing the distribution of virtual machines per hosts.**

- 5 Click **Presentation** and select the **Distribution** view type.

The view type is the way the information is displayed.

- a From the **Visualization** drop-down menu, select **Pie Chart**.
- b From the Distribution Type configurations, select **Discrete distribution**.

Leave **Max number of buckets** deselected because you do not know the number of hosts on each vCenter Server instance. If you specify a number of buckets and the hosts are more than that number, one of the slices shows unspecified information labeled Others.

- 6 Click **Subjects** to select the object type that applies to the view.

- a From the drop-down menu, select **Host System**.

The Distribution view is visible at the object containers of the subjects that you specify during the view configuration.

- 7 Click **Data** and in the filter text box enter **Total Number of VMs**.

- 8 Select **Summary > Total Number of VMs** and double-click to add the metric.

- 9 Retain the default metric configurations and click **Save**.

Run a vRealize Operations Manager View

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Environment**.
- 2 In the left pane, navigate to a vCenter Server instance and click the **Details** tab.
All listed views are applicable for the vCenter Server instance.
- 3 From the **All Filters** drop-down menu on the left, select **Type > Distribution**.
You filter the views list to show only distribution type views.
- 4 Navigate to and click the **Virtual Machines Distribution** view.

The bottom pane shows the distribution view with information about this vCenter Server. Each slice represents a host and the numbers on the far left show the number of virtual machines.

Export a vRealize Operations Manager View

To use a view in another vRealize Operations Manager, you export a content definition XML file.

If the exported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Click the gear icon and select **Export View**.
- 3 In the list of views, navigate to and click the **Virtual Machines Distribution** view .
- 4 Select a location on your local system to save the XML file and click **Save**.

Import a vRealize Operations Manager View

To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Click the gear icon and select **Import View**.
- 3 Browse to select the Virtual Machines Distribution content definition XML file and click **Import**.

If the imported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

Note The imported view overwrites if a view with the same name exists. All report templates that use the existing view are updated with the imported view.

Reports

A report is a scheduled snapshot of views and dashboards. You can create it to represent objects and metrics. It can contain table of contents, cover page, and footer.

With the vRealize Operations Manager reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.



Create Reports

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_reports_vrops)

Report Templates Tab

On the **Report Templates** tab you can create, edit, delete, clone, run, schedule, export, and import templates.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Report Templates** to access the Reports Templates tab.

All templates that are applicable for the selected object are listed on the **Report Templates** tab. You can order them by report name, subject, date they were modified, last run, or owner.

You can filter the templates list by adding a filter from the right side of the panel.

Table 7-20. Predefined Filter Groups

Filter Group	Description
Name	Filter by the template name. For example, you can list all reports that contain <i>my template</i> in their name by typing my template .
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by those objects.

vSphere users must be logged in until the report generation is complete. If you log out or your session expires, the report generation fails.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

Generated Reports Tab

All reports that are generated for a selected object are listed on the **Generated Reports** tab.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Generated Reports** to access the Generated Reports tab.

You can order the reports by the date and time that they were created, the report name, the owner, or their status. If the report is generated through a schedule, the owner is the user who created the schedule.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

You can filter the reports list by adding a filter from the right side of the panel.

Table 7-21. Predefined Filter Groups

Filter Group	Description
Report Name	Filter by the report template name. For example, you can list all reports that contain <i>my template</i> in their name by typing my template .
Template	Filter by the report template. You can select a template from a list of templates applicable for this object.
Completion Date/Time	Filter by the date, time, or time range.
Status	Filter by the status of the report. On each data node, only one report can be processed. Therefore, reports that are queued can be moved to the processed state only after the previous report on the specific node has failed or completed. The maximum queue time is restricted to 4 hours. After 4 hours, if processing of the report has not started, the report is marked as failed.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by those objects.

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

Create and Modify a Report Template

You create a report to generate a scheduled snapshot of views and dashboards. You can track current resources and predict potential risks to the environment. You can schedule automated reports at regular intervals.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Reports**.
- 2 On the **Report Templates** tab, click the **New Template** icon to create a template.
- 3 Complete the steps in the left pane to:
 - a Enter a name and description for the report template.
[Name and Description Details](#)
 - b Add a view or a dashboard.
[Views and Dashboards Details](#)
 - c Select an output for the report.
[Formats Details](#)
 - d Select the layout options.
[Layout Options Details](#)
- 4 Click **Save**.

5 From the Report Templates tab, click **Edit Template** to modify the report template.

Name and Description Details

The name and description of the report template as they appear in the list of templates on the **Report Templates** tab.

Where You Add Name and Description

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Name and Description**.

Table 7-22. Name and Description Options in the Report Template Workspace

Option	Description
Name	Name of the template as it appears on the Report Templates tab.
Description	Description of the template.

Views and Dashboards Details

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs.

Where You Add Views and Dashboards

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Views and Dashboards**. If you create a template, complete the required previous steps of the workspace.

How You Add Views and Dashboards

To add a view or a dashboard to your report template, select it from the list on the left pane and drag it to the main panel. You can drag the views and dashboards in the main panel to reorder them. You can select a portrait or landscape orientation for each view or dashboard from the drop-down menu next to its title.

Table 7-23. Views and Dashboards Options in the Report Template Workspace

Option	Description
Data type	Select Views or Dashboards to display a list of available views or dashboards that you can add to the template.
Create View	Create a view directly from the template workspace. This option is available when you select Views from the Data type drop-down menu.
Edit View	Edit a view directly from the template workspace. This option is available when you select Views from the Data type drop-down menu.
Create Dashboard	Create a dashboard directly from the template workspace. This option is available when you select Dashboards from the Data type drop-down menu.
Edit Dashboard	Edit a dashboard directly from the template workspace. This option is available when you select Dashboards from the Data type drop-down menu.
Search	Search for views or dashboards by name. To see the complete list of views or dashboards, delete the search box contents and press Enter.
List of views	List of the views that you can add to the template. This list is available when you select Views from the Data type drop-down menu.
List of dashboards	List of the dashboards that you can add to the template. This list is available when you select Dashboards from the Data type drop-down menu.
Preview of views and dashboards	In the main panel, you see a preview of the views and dashboards that you add. When you create a template in the context of an object from the environment, you see a live preview of the views and dashboards.

Formats Details

The formats are the outputs in which you can generate the report.

Where You Add Formats

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Formats** to select a format for the report template. If you create a template, complete the required previous steps of the workspace.

Table 7-24. Formats Options in the Report Template Workspace

Option	Description
PDF	With the PDF format, you can read the reports, either on or off line. This format provides a page-by-page view of the reports, as they appear in printed form.
CSV	In the CSV format, the data is in a structured table of lists.

Layout Options Details

The report template can contain layout options such as a cover page, table of contents, and footer.

Where You Add Layout Options

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Layout Options**. If you create a template, complete the required previous steps of the template.

Table 7-25. Layout Options in the Report Template Workspace

Option	Description
Cover Page	Can contain an image up to 5 MB. The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page.
Table of contents	Provides a list of the template parts, organized in the order of their appearance in the report.
Footer	Includes the date when the report is created, a note that the report is created by VMware vRealize Operations Manager, and page number.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports SMB version 2.0.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.
- 2 From the toolbar, click the **Add** icon.

- 3 From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

The dialog box expands to include your plug-in instance settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.
Password	The password for the domain user account.
Network share root	<p>The path to the root folder where you want to save the reports. You can specify subfolders for each report when you configure the schedule publication.</p> <p>You must enter an IP address. For example, <code>\\IP_address\ShareRoot</code>. You can use the host name instead of the IP address if the host name is resolved to an IPv4 when accessed from the vRealize Operations Manager host.</p> <p>Note Verify that the root destination folder exists. If the folder is missing, the Network Share plug-in logs an error after 5 unsuccessful attempts.</p>

- 6 Click **Test** to verify the specified paths, credentials, and permissions.

The test might take up to a minute.

- 7 Click **Save**.

The outbound service for this plug-in starts automatically.

- 8 (Optional) To stop an outbound service, select an instance and click **Disable** on the toolbar.

Results

This instance of the Network Share plug-in is configured and running.

What to do next

Create a report schedule and configure it to send reports to your shared folder.

Configuring Application Monitoring with Wavefront

8

You can monitor application services supported by VMware Application Proxy in Wavefront. You can also manage the life cycle of agents and application services on virtual machines.

For example, as an administrator, you might need to ensure that the infrastructure provided for running the application services is sufficient and that there are no problems. If you receive a complaint that a particular application service is not working properly or is slow, you can troubleshoot by looking at the infrastructure on which the application is deployed. You can view important metrics related to the applications and share the information with the team managing the applications. You can use vRealize Operations Manager to deploy the agents and send the related application data to Wavefront. You can view the data in the relevant Wavefront dashboard and share it with the team so that they can troubleshoot the application service.

To monitor your application services and collect metrics for applications services supported by VMware Application Proxy, follow these steps in vRealize Operations Manager:

- 1 Provide configuration details to activate your Wavefront account. If you do not have a Wavefront account, you can create a free 30-day trial Wavefront account.

For more information, see [Configure a Wavefront Account](#).

- 2 Download the VMware Application Proxy by clicking **Download** in the **Manage VMware Application Proxy** panel.

For information about deploying VMware Application Proxy, see [Deploy VMware Application Proxy](#).

- 3 Configure an application proxy.

For information about configuring VMware Application Proxy, see [Manage VMware Application Proxy](#) and [Add and Configure an Application Proxy](#).

- 4 Install agents on selected VMs and discover and manage application services.

For more information, see [Manage Agents in Virtual Machines](#).

- 5 Monitor the metrics in Wavefront.

For more information, see [Monitor Metrics in Wavefront](#) and the [Wavefront](#) documentation.

This chapter includes the following topics:

- [Configure a Wavefront Account](#)

- [Manage VMware Application Proxy](#)
- [Manage Agents in Virtual Machines](#)
- [Monitor Metrics in Wavefront](#)

Configure a Wavefront Account

Use the **Configure a VMware Wavefront Account** panel to configure a Wavefront account to monitor and collect metrics of applications supported by VMware Application Proxy.

You can also create a 30-day free trial account from the **Configure a VMware Wavefront Account** panel. For more details, see [Create a Wavefront Trial Account](#).

Where You Configure the Wavefront Account

To configure the Wavefront account, in the menu, select **Home**, and then from the left pane select the **Application Monitoring (Wavefront)**. Click the **Configure a VMware Wavefront Account** panel in the right pane and select **I have an Existing Account**.

The top left corner of the panel displays the status of the operation.

Table 8-1. Status Details

Status	Description
Orange warning symbol	Indicates that you have not configured a Wavefront account.
Green tick mark	Indicates that you have configured a Wavefront account. Refresh the screen after you have configured the Wavefront account.

Configure the Wavefront Account

You must have a Wavefront account. If you do not have a Wavefront account, you can sign up for a 30-day trial by clicking **Create New Account** from the **Configure a VMware Wavefront Account** panel.

Procedure

- 1 Enter your Wavefront service URL, for example, `http://longboard.wavefront.com`.
- 2 Enter your API token for the Wavefront account.
You receive the Wavefront URL and the API token in an email.
- 3 Click **Test Connection** to validate the connection.
- 4 You can either:
 - Click **Save and Proceed** to complete configuring your Wavefront account.
 - Click **Delete** to delete Wavefront configuration details after you have saved the details.
 - Click **Cancel**. You can configure or create a Wavefront account.

Create a Wavefront Trial Account

If you do not have a Wavefront account, you can create a free 30-day trial account.

Procedure

- 1 Click **Create New Account** in the **Configure a VMware Wavefront Account** panel in vRealize Operations Manager.

You are redirected to the Wavefront free trial website.

- 2 Enter your first name, last name, work email, and company name in the text boxes provided.
- 3 Select the size of your company and your country of residence.
- 4 Select **Sign Up Now**.

Results

You receive an email with the Wavefront URL and the API token. Use these details to configure the Wavefront account in vRealize Operations Manager. For more information about configuring the Wavefront account, see [Configure the Wavefront Account](#).

Manage VMware Application Proxy

The Application proxies you add and configure are displayed in the **Manage VMware Application Proxy** panel.

You can view the name of the VMware Application Proxy added and the number of vCenters managed, in the **Manage VMware Application Proxy** panel.

If there are newer versions of VMware Application Proxy available, you see the **Updates available** label on the top right corner of the **Manage VMware Application Proxy** panel. Point to the **Updates available** label to view a tooltip that displays a URL. You can upgrade VMware Application Proxy from this URL.

Where You Can Configure the Application Proxy

To configure an application proxy, from the menu, select **Home**, and then from the left pane select **Application Monitoring (Wavefront)**. From the right pane, click the **Manage VMware Application Proxy** panel.

Table 8-2. Options

Options	Description
Download	You can download VMware Application Proxy. For information about deploying VMware Application Proxy, see Deploy VMware Application Proxy .
Add	<p>You can map a vCenter Server with a VMware Application Proxy.</p> <p>When you click Test Connection to validate the connection, the Review and Accept Certificate dialog box is displayed. Click Accept if you trust the certificate.</p>

Table 8-2. Options (continued)

Options	Description
Edit	<p>You can modify the VMware Application Proxy configuration details or the details of the vCenter Servers that are managed.</p> <p>After you modify the details and click Test Connection, the Review and Accept Certificate dialog box is displayed if you have not already accepted the certificate. Click Accept if you trust the certificate. The connection is then validated.</p>
Delete	<p>You can delete the application proxy. Data will not be sent to Wavefront. Ensure that you uninstall the agents from the VMs that are monitored before you delete the application proxy.</p>

You can also view specific details from the options in the data grid.

Table 8-3. Data Grid Options

Option	Description
Name	Displays the FQDN of the VMware Application Proxy.
Version	Displays the version of VMware Application Proxy. A gray dot is displayed if there is a newer version of VMware Application Proxy available.
vCenters Managed	Displays the number of vCenter Servers mapped to the VMware Application Proxy.
Proxy Server Status	<p>Indicates the health of the VMware Application Proxy.</p> <ul style="list-style-type: none"> ■ Green. Indicates that the VMware Application Proxy is healthy. ■ Red. Indicates that the VMware Application Proxy is not healthy. <p>Point to this cell to view a tooltip that displays the cause if the health status is red.</p> <p>The progress status is displayed when data collection has not started.</p>
Wavefront Connection Status	<p>Indicates the health of the VMware Application Proxy's connection to Wavefront.</p> <ul style="list-style-type: none"> ■ Green. Indicates a healthy connection. ■ Red. Indicates that the connection is not healthy. <p>Point to this cell to view a tooltip that displays the cause if the health status is red.</p> <p>The progress status is displayed when data collection has not started.</p>

The top left corner of the Manage VMware Application Proxy panel displays the status of the operation.

Table 8-4. Status Details

Status	Description
Orange warning symbol	Indicates that you have not added VMware Application Proxy.
Green tick mark	Indicates that all instances of VMware Application Proxy are healthy.

Deploy VMware Application Proxy

Use a vSphere client to deploy VMware Application Proxy. You can deploy the VMware Application Proxy OVA template from a URL or from a file.

Prerequisites

You must have the URL to the VMware Application Proxy OVA template before you proceed. Alternately, you can download the VMware Application Proxy OVA file after you log in to vRealize Operations Manager. Navigate to the **Application Monitoring with VMware Wavefront** page from the **Quick Start** page and click the download link in step 2.

For critical time sourcing, use the Network Time Protocol (NTP). You must ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts and vRealize Operations Manager.

Procedure

- 1 Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select **Deploy OVF Template**.
The **Deploy OVF Template** wizard opens.
- 2 Select **Deploy OVF Template**.
The **Deploy OVF Template** wizard opens.
- 3 On the **Deploy OVF template** page do one of the following and click **Next**:
 - ◆ If you have a URL to the OVA template which is located on the Internet, type the URL in the URL field. Supported URL sources are HTTP and HTTPS
 - ◆ If you have downloaded the VMware Application Proxy OVA file, click **Local file** and browse to the location of the file and select it.
- 4 On the **Select a name and folder** page, enter a unique name for the virtual machine or vAPP, select a deployment location, and click **Next**.

The default name for the virtual machine is the same as the name of the selected OVF or OVA template. If you change the default name, choose a name that is unique within each vCenter Server virtual machine folder.

The default deployment location for the virtual machine is the inventory object where you started the wizard.

- 5 On the **Select a resource** page, select a resource where to run the deployed VM template, and click **Next**.
- 6 On the **Review details** page, verify the OVF or OVA template details and click **Next**.

Option	Description
Product	VMware Application Proxy.
Version	Version number of the VMware Application Proxy.
Vendor	VMWare.
Publisher	Publisher of the OVF or OVA template, if a certificate included in the OVF or OVA template file specifies a publisher.
Download size	Size of the OVF or OVA file.
Size on disk	Size on disk after you deploy the OVF or OVA template.

- 7 On the **Accept license agreements** page, click **Accept** and then **Next**.
- 8 In the **Select configuration** page, select the size of the deployment.
- 9 On the **Select storage** page, define where and how to store the files for the deployed OVF or OVA template.
 - a Select a VM Storage Policy.
This option is available only if storage policies are enabled on the destination resource.
 - b (Optional) Enable the **Show datastores from Storage DRS clusters** check box to choose individual datastores from Storage DRS clusters for the initial placement of the virtual machine.
 - c Select a datastore to store the deployed OVF or OVA template.
The configuration file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine or vApp and all associated virtual disk files.
- 10 On the **Select networks** page, select a source network and map it to a destination network. Click **Next**. The source network must have a static FQDN name or static DNS.
The Source Network column lists all networks that are defined in the OVF or OVA template.
- 11 In the Customize template page, provide inputs to configure the VMware Application Proxy deployment.

Configuration	Description
API Admin User's Password	Enter a password for the VMware Application Proxy API admin. The username is admin@ucp.local. This password should be used when configuring this Application Proxy in vRealize Operations Manager.
Networking Properties	Verify the networking properties.

- 12 On the **Ready to complete** page, review the page and click **Finish**.

- 13** After the OVA deployment is complete, you can log in to the virtual appliance from vCenter Server. Right click the virtual appliance that you installed. Click **Open Console**. Use the following credentials to log in:

Log In Details	Value
Username	root
Password	vmware

- 14** Change the root user password.
- 15** Enable the sshd service to access the virtual machine through ssh.

What to do next

- Perform the post-installation tasks.
- Log in to vRealize Operations Manager and configure the agents to connect to Wavefront.

Configure Network Time Protocol Settings

After you install or upgrade to VMware Application Proxy version 1.x.x, you must set up accurate timekeeping as part of the deployment. If the time settings between VMware Application Proxy and vRealize Operations Manager are not synchronized, you will face agent installation and metric collection issues. Ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts and vRealize Operations Manager using the Network Time Protocol (NTP).

Procedure

- 1** Log in to the VMware Application Proxy VM and modify the `ntp.conf` file available in `/etc/ntp.conf` by adding following in the following format:

```
server time.vmware.com
```

Note Replace `time.vmware.com` with a suitable time server setting. You can use the FQDN or IP of the time server.

- 2** Enter the following command to start the NTP daemon:

```
systemctl start ntpd
```

- 3** Enter the following command to enable the NTP daemon:

```
systemctl enable ntpd
```

- 4** Run the following command to verify if NTP is configured correctly:

```
ntpstat
```

If NTP is synchronized correctly, you will see a message similar to the following:

```
synchronised to NTP server (10.113.60.176) at stratum 3

time correct to within 50 ms

polling server every 64 s
```

Add and Configure an Application Proxy

You can add and configure application proxies from the **Manage VMware Application Proxy** panel to manage the life cycle of agents and application services.

To add and configure a VMware Application Proxy, in the menu, click **Home**, and then in the left pane select **Application Monitoring (Wavefront)**. From the right pane, click the **Manage VMware Application Proxy** panel.

Note Time synchronization between VMware Application Proxy and vRealize Operations Manager is mandatory when you add an application proxy. If the time settings are not synchronized, you face problems such as, a failed test connection when you add an application proxy, agent installation issues, and issues in metrics collection after the agent is installed. Eventually, metrics are not displayed in Wavefront dashboards. For more information, see the called **Troubleshoot Agent Installation and Metric Collection Issues** in the [VMware Application Proxy Guide](#) on Solution Exchange.

For more troubleshooting information on VMware Application Proxy, see the chapter called **Troubleshooting your Deployment** in the [VMware Application Proxy Guide](#).

Prerequisites

- Verify that you have configured a vCenter adapter.
- Download and deploy VMware Application Proxy.

You can download **VMware Application Proxy** by clicking **Download** in the **Manage VMware Application Proxy** page.

For information about deploying the VMware Application Proxy, see [Deploy VMware Application Proxy](#).

Procedure

- 1 To configure a VMware Application Proxy, click **Add** from the **Manage VMware Application Proxy** panel.
- 2 In the **Configure Application Proxy** page, enter the following details:
 - a FQDN of the VMware Application Proxy you have configured during the installation of VMware Application Proxy.
 - b You cannot modify the user name which is **admin**.

- c The API password of the VMware Application Proxy you have configured during the installation of VMware Application Proxy.
- d Click **Next**.

3 From the **Map vCenters** page, complete the following steps:

- a Select the vCenter Servers to which you want to map the VMware Application Proxy.
If you have mapped a vCenter Server to a VMware Application Proxy, it is not displayed in the drop-down menu.
- b The vCenter Servers that are mapped to the VMware Application Proxy are displayed on the page.
- c Click **Test Connection** to validate the connection. The **Review and Accept Certificate** dialog box is displayed. Click **Accept** if you trust the certificate.

If the mapped vCenter Server turns red, it signifies that vRealize Operations Manager cannot communicate with the VMware Application Proxy. If the mapped vCenter Server turns green, it signifies that vRealize Operations Manager can communicate with the VMware Application Proxy.
- d Click **Next**.

4 From the **Summary** page, you view details such as the FQDN, user name, and the vCenter Servers that are mapped to an instance of the VMware Application Proxy.

It might take up to 5 minutes to get the status of VMware Application Proxy.

- a Click **Finish**.

What to do next

Install agents on the VMs you prefer and manage the application services.

Manage Agents in Virtual Machines

After you have configured the VMware Application Proxy and mapped it to a vCenter Server, you can manage the agents on the VMs from the **Manage Agents in Virtual Machines** panel. You can view the data centers, hosts, and clusters available in the vCenter Servers you have mapped to VMware Application Proxy. You can install, uninstall, start, stop, and update the agents on the VMs. You can also discover and manage the services on each agent that you install.

Where You Manage the Agents

To manage the agents and application services, in the menu, select **Home**, and then from the left pane select **Application Monitoring (Wavefront)**. From the right pane, click the **Manage Agents in Virtual Machines** panel.

Table 8-5. Options

Options	Description
Virtual Machine Filter	<ul style="list-style-type: none"> ■ Lists the data centers, hosts, clusters, and VMFolders available in the vCenter Servers you have mapped to VMware Application Proxy. ■ Lists the vRealize Automation objects such as tenant, business group, blueprint, and deployment as filters. You can use these filters if you use vRealize Automation for virtual machine provisioning and application deployment and if you use vRealize Operations Manager to monitor the vRealize Automation infrastructure. <p>To view vRealize Automation objects as filters, complete the following prerequisites:</p> <ul style="list-style-type: none"> ■ Deploy VMware Application Proxy. ■ Configure the vCenter adapter that you added as an endpoint for provisioning in vRealize Automation. ■ Configure the vRealize Automation adapter instance in vRealize Operations Manager. ■ Wait for the resources to be discovered. <p>You can filter the VMs for each object that you select.</p>
Manage Agent	<p>Allows you to install, uninstall, start, stop, and update the agent.</p> <ul style="list-style-type: none"> ■ Install. Installs the agents on the selected VM. Select the VMs on which you want to install the agent and select Manage Agent > Install. ■ Uninstall. Uninstalls the agent. Select the VMs on which you want to uninstall the agent and select Manage Agent > Uninstall. ■ Start. If you have temporarily stopped sending metrics to Wavefront, you can use this option to start data collection for the application service. ■ Stop. During a maintenance period, you can temporarily stop sending application service metrics to Wavefront. Select the VMs on which you want to stop the agent and select Manage Agent > Stop. ■ Update. Updates agents that are at a lower version. Select the VMs on which you want to update the agent and select Manage Agent > Update. After the agents are updated, the agent status changes to Update Success.
Manage Service	You can manage the application services that are discovered on the virtual machines where the agents are installed.
All Filters	Filters the VMs based on the name of the VM, the operating system it runs on, the application service discovered, and the power status of the VM.

You can also view specific details from the options in the data grid.

Table 8-6. Data Grid Options

Option	Description
VM Name	Name of the virtual machine.
Version	Version of the VMware Application Proxy agent on the VM. A gray dot is displayed if the VM requires an update.
vCenter Name	Name of the vCenter Adapter instance to which that VM resource belongs.
Operating System	Operating system installed on the VM.

Table 8-6. Data Grid Options (continued)

Option	Description
Agent Status	<p>Displays the status of the agent at the end point.</p> <ul style="list-style-type: none"> ■ Blue icon. Indicates that the agent is not installed. ■ Green dot. Indicates that the agent is running. ■ Red dot. Indicates that the agent has stopped.
Last Operation Status	<p>Status of the last operation. The possible values are:</p> <ul style="list-style-type: none"> ■ No Operation ■ Install Success ■ Install Failed ■ Install In Progress ■ Start Success ■ Start Failed ■ Start In Progress ■ Stop Success ■ Stop Failed ■ Stop In Progress ■ Update Success ■ Update Failed ■ Update In Progress ■ Uninstall Success ■ Uninstall Failed ■ Uninstall In Progress
VM State	<p>Power status of the VMs. The possible values are:</p> <ul style="list-style-type: none"> ■ Powered On ■ Powered Off
Services Discovered	<p>List of the supported application services discovered on the VM.</p> <ul style="list-style-type: none"> ■ A red dot against the application service indicates that the application service has been activated but there is a problem with data collection. ■ A gray dot before the application service indicates that the agent requires reactivation. The application service must be reactivated. For reactivation, see Activate and Deactivate an Application Service for more information. ■ A steel blue dot indicates that the agents have stopped. ■ A green dot against the application service indicates that the application service is activated. ■ If an application service has been deactivated or not activated, you will not see a symbol displayed against the application service. ■ After you have added the parameters and activated the application service, the progress status is displayed until data collection starts.

The top left corner of the Manage Agents in Virtual Machines panel displays the status of the operation.

Table 8-7. Status Details

Status	Description
Orange warning symbol	Indicates that none of the VMs are bootstrapped.
Green tick mark	Indicates that at least one of the VMs is bootstrapped and the agent is running.

To manage the agent, follow these steps:

- 1 Install the agent.
For more information, see [Install an Agent](#).
- 2 Manage the application services on each agent.
For more information, see [Manage Application Services](#).
- 3 Stop and start the agents on the VMs.
- 4 Uninstall the agent.
For more information, see [Uninstall an Agent](#).
- 5 Update agents that are at a lower version.

Install an Agent

You must select the VMs on which you want to install the agent. If you have upgraded an existing installation of VMware Application Proxy, reinstall the agents that you have previously installed.

Prerequisites

- Time synchronization between VMware Application Proxy, vRealize Operations Manager, ESX hosts, and Windows and Linux target VMs is mandatory for secure communication.
- Account Privilege Prerequisites
 - Windows end points. To install agents,
 - The user must be either an administrator, or
 - A non-administrator who belongs to the administrator group with UAC disabled on the operating system.
To disable UAC (previously known as LUA) on Windows, complete the following step:
 - In the registry path HKLM:\Software\Microsoft\Windows\CurrentVersion\Policies\System, set the value for the key EnableLUA to 0.
 - Reboot the machine for the changes to take effect.
 - Linux end points. To install agents,
 - The user must be a root user, or
 - A non-root user with the following capability:
 - Password-less sudo elevation access for a non-root user or a non-root user group.

To enable password-less sudo elevation access for a user called NRU, add NRU ALL=(ALL:ALL) NOPASSWD: ALL to /etc/sudoers.

To enable password-less sudo elevation access for a user group called NRUG, add %NRUG ALL=(ALL:ALL) NOPASSWD: ALL to /etc/sudoers.

Procedure

- 1 From the **Manage Agents in Virtual Machines** panel, select **Manage Agent > Install**. You see the **Manage Agent** dialog box.
- 2 From the **How do you want to provide VM Credentials** page, complete the following steps:
 - a If you have a common user name and password for all the VMs, select the **Common username and password** option.
 - b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
 - c Click **Next**.
- 3 From the **Provide Credentials** page, depending on whether you have a common credential for all VMs or different credentials for all VMs, enter the following details:
 - a If the selected VMs have a common user name and password, enter the common user name and password.
 - b For different user names and passwords for each VM, download the CSV template and add the required details such as the user name, password for each VM. Use the **Browse** button to select the template.
 - c Click **Next**.
- 4 From the **Summary** page, you can view the list of VMs on which the agent will be deployed.
- 5 Click **Deploy Agent**. Refresh the UI to view the agents that are installed.

The agent discovers the application services that are installed on the VMs and the application services are displayed in the **Services Discovered** column in the **Manage Agents in Virtual Machine** panel. You can view the status of agent installation from the **Agent Status** column in the **Manage Agents in Virtual Machine** panel.

What to do next

You can manage the services on each agent.

Manage Application Services

You can manage the application services supported by VMware Application Proxy on the VMs where the agents are installed.

Procedure

- 1 Select a VM on which the agent has been installed and the application services have been discovered, from the **Manage Agents in Virtual Machines** panel.
- 2 Select **Manage Service** and then from the drop-down menu select the **service name**. You see the **Plugin Activation** dialog box.
- 3 By default, all metrics are collected for the activated application service.
- 4 Activate data collection for the application service.
- 5 Enter the relevant settings for the application service.
- 6 Click **Confirm**.

Fields with a star are mandatory.

For more information about the status details that appear against the application services in the Services Discovered column, see the table called Data Grid Options in [Manage Agents in Virtual Machines](#).

For information about supported application services and their properties, see *Supported Application Services* in the *VMware Application Proxy Guide* on Solution Exchange.

What to do next

You can view the metrics collected for each application service in the Wavefront dashboards.

Activate and Deactivate an Application Service

To monitor application services running on the target VMs, VMware Application Proxy plugins must be configured in the target VMs after the agent is installed.

After you have installed the agent, you can choose to activate or deactivate VMware Application Proxy plugins to monitor application services. You can also reactivate plugins that need to be monitored.

Activate an Application Service

To monitor an application service, complete the following steps:

- 1 Navigate to the **Manage Agents in Virtual Machines** panel.
- 2 Select the VM on which agent is already installed.
- 3 Select **Manage Service** and then from the drop-down menu select the **service name**.
- 4 Activate the application service from the **PluginActivation** dialog box.
- 5 Enter the password and click **Confirm**.

For more information about the status details that appear against the application services in the Services Discovered column, see the table called Data Grid Options in [Manage Agents in Virtual Machines](#).

The following special characters are permitted in the DB user field: ' [] {} () , . < > ? : ! | / ~ @ # \$ % ^ & * - _ +=

You can provide DB name lists in the following format ['DBNAME_1', 'DBNAME_2', 'DBNAME_3'] where DBNAME_1, DBNAME_2, DBNAME_3 must not contain quotes such as ' and ''.

Note You can activate one application service at a time. When multiple VMs are selected, the **Manage Service** option is disabled.

Deactivate an Application Service

To deactivate a plugin to stop monitoring the application service that is sending data to Wavefront, complete the following steps:

- 1 Navigate to the **Manage Agents in Virtual Machines** panel.
- 2 Select the VM on which the agent is already installed.
- 3 Select **Manage Service** and then from the drop-down menu select the **service name**.
- 4 Deactivate the application service from the **PluginActivation** dialog box and click **Confirm**.

Uninstall an Agent

You must select the VMs on which you want to uninstall the agent.

Procedure

- 1 From the **Manage Agents in Virtual Machines** panel, select **Manage Agent > Uninstall**. You see the **Manage Agent** dialog box.
- 2 From the **How do you want to provide VM Credentials** page, complete the following steps:
 - a If you have a common user name and password for all the VMs, select the **Common username and password** option.
 - b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
 - c Click **Next**.
- 3 From the **Provide Credentials** page, depending on whether you have a common credential for all VMs or different credentials for all VMs, enter the following details:
 - a If your VM has a single user name and password, enter the common user name and password.
 - b For multiple user names and passwords for each VM, download the CSV template and add the details. Use the **Browse** button to select the template.
 - c Click **Next**.
- 4 From the **Summary** page, you can view the list of VMs on which the agent is deployed.

- 5 Click **Remove Agent**. Refresh the UI to view the progress of agent uninstallation.

The **Agent Status** and **Services Discovered** columns in the workspace indicate that uninstallation is complete and that there are no application services discovered on each agent.

Monitor Metrics in Wavefront

To monitor metrics for the application services you have activated, open Wavefront and view the dashboards that are populated with data.

You can also access Wavefront by selecting **Home** in the menu, and then from the left pane, select **Application Monitoring (Wavefront)**. From the right pane, click the **View in Wavefront** button.

For VMware Application Proxy 1.0

From the **Wavefront** home page, select **Integrations** and click the application service you have activated. From the **Dashboard** tab, click the application link. Select **View > Source** and enter the <ID of the vCenter server> and the <unique VM number> in the search field in the format <VCID>_<VMMOR>. For example, 0e6304a2-df1f-4043-ae40-0d891d443387_vm-99.

For VMware Application Proxy 1.0.0.1 and 1.1

From the **Wavefront** home page, select **Integrations** and click the application service you have activated. From the **Dashboard** tab, click the application link. Select **View > Source** and enter the Virtual Machine name.

If you have created custom dashboards, you must recreate the dashboards as the Virtual Machine Name is now the source tag and there is no separate vm_name point tag in VMware Application Proxy 1.0.0.1. However, there is a new point tag for all metrics in VMware Application Proxy 1.1. This is the *vc_uuid* tag that carries the UUID of the vCenter Server that manages the relevant virtual machine.

For more information, see the [Wavefront](#) documentation.

Configuring Administration Settings

9

After vRealize Operations Manager is installed and configured, you can use administration settings to manage your environment. You find most administration settings under the Administration selection of the vRealize Operations Manager interface.

This chapter includes the following topics:

- [Managing Users and Access Control in vRealize Operations Manager](#)
- [vRealize Operations Manager Passwords and Certificates](#)
- [Modifying Global Settings](#)
- [Transfer Ownership of Dashboards and Report Schedules](#)
- [Create a vRealize Operations Manager Support Bundle](#)
- [Customizing Icons](#)

Managing Users and Access Control in vRealize Operations Manager

To ensure security of the objects in your vRealize Operations Manager instance, as a system administrator you can manage all aspects of user access control. You create user accounts, assign each user to be a member of one or more user groups, and assign roles to each user or user group to set their privileges.

Users must have privileges to access specific features in the vRealize Operations Manager user interface. Access control is defined by assigning privileges to both users and objects. You can assign one or more roles to users, and enable them to perform a range of different actions on the same types of objects. For example, you can assign a user with the privileges to delete a virtual machine, and assign the same user with read-only privileges for another virtual machine.

User Access Control

You can authenticate users in vRealize Operations Manager in several ways.

- Create local user accounts in vRealize Operations Manager.

- Use VMware vCenter Server® users. After the vCenter Server is registered with vRealize Operations Manager, configure the vCenter Server user options in the vRealize Operations Manager global settings to enable a vCenter Server user to log in to vRealize Operations Manager. When logged into vRealize Operations Manager, vCenter Server users access objects according to their vCenter Server-assigned permissions.
- Add an authentication source to authenticate imported users and user group information that resides on another machine.
 - Use LDAP to import users or user groups from an LDAP server. LDAP users can use their LDAP credentials to log in to vRealize Operations Manager.
 - Create a single sign-on source and import users and user groups from a single sign-on server. Single sign-on users can use their single sign-on credentials to log in to vRealize Operations Manager and vCenter Server. You can also use Active Directory through single sign-on by configuring the Active Directory through single sign-on and adding the single sign-on source to vRealize Operations Manager.

User Preferences

To determine the display options for vRealize Operations Manager, such as colors for the display and health chart, the number of metrics and groups to display, and whether to synchronize system time with the host machine, you configure the user preferences on the top toolbar.

Users of vRealize Operations Manager

Each user has an account to authenticate them when they log in to vRealize Operations Manager.

The accounts of local users and LDAP users are visible in the vRealize Operations Manager user interface when they are set up. The accounts of vCenter Server and single sign-on users only appear in the user interface after a user logs in for the first time. Each user can be assigned one or more roles, and can be an authenticated member of one or more user groups.

Local Users in vRealize Operations Manager

When you create user accounts in a local vRealize Operations Manager instance, vRealize Operations Manager stores the credentials for those accounts in its global database, and authenticates the account user locally.

Each user account must have a unique identity, and can include any associated user preferences.

If you are logging in to vRealize Operations Manager as a local user, and on occasion receive an **invalid password** message, try the following workaround. In the Login page, change the Authentication Source to **All vCenter Servers**, change it back to **Local Users**, and log in again.

vCenter Server Users in vRealize Operations Manager

vRealize Operations Manager supports vCenter Server users. To log in to vRealize Operations Manager, vCenter Server users must be valid users in vCenter Server.

Roles and Associations

A vCenter Server user must have either the vCenter Server Admin role or one of the vRealize Operations Manager privileges, such as PowerUser which assigned at the root level in vCenter Server, to log in to vRealize Operations Manager. vRealize Operations Manager uses only the vCenter privileges, meaning the vRealize Operations Manager roles, at the root level, and applies them to all the objects to which the user has access. After logging in, vCenter Server users can view all the objects in vRealize Operations Manager that they can already view in vCenter Server.

Logging in to vCenter Server Instances and Accessing Objects

vCenter Server users can access either a single vCenter Server instance or multiple vCenter Server instances, depending on the authentication source they select when they log in to vRealize Operations Manager.

- If users select a single vCenter Server instance as the authentication source, they have permission to access the objects in that vCenter Server instance. After the user has logged in, an account is created in vRealize Operations Manager with the specific vCenter Server instance serving as the authentication source.
- If users select **All vCenter Servers** as the authentication source, and they have identical credentials for each vCenter Server in the environment, they see all the objects in all the vCenter Server instances. Only users that have been authenticated by all the vCenter Servers in the environment can log in. After a user has logged in, an account is created in vRealize Operations Manager with all vCenter Server instances serving as the authentication source.

vRealize Operations Manager does not support linked vCenter Server instances. Instead, you must configure the vCenter Server adapter for each vCenter Server instance, and register each vCenter Server instance to vRealize Operations Manager.

Only objects from a specific vCenter Server instance appear in vRealize Operations Manager. If a vCenter Server instance has other linked vCenter Server instances, the data does not appear.

vCenter Server Roles and Privileges

You cannot view or edit vCenter Server roles or privileges in vRealize Operations Manager. vRealize Operations Manager sends roles as privileges to vCenter Server as part of the vCenter Server Global privilege group. A vCenter Server administrator must assign vRealize Operations Manager roles to users in vCenter Server.

vRealize Operations Manager privileges in vCenter Server have the role appended to the name. For example, vRealize Operations Manager ContentAdmin Role, or vRealize Operations Manager PowerUser Role.

Read-Only Principal

A vCenter Server user is a read-only principal in vRealize Operations Manager, which means that you cannot change the role, group, or objects associated with the role in vRealize Operations Manager. Instead, you must change them in the vCenter Server instance. The role applied to the root folder applies to all the objects in vCenter Server to which a user has privileges. vRealize

Operations Manager does not apply individual roles on objects. For example, if a user has the PowerUser role to access the vCenter Server root folder, but has read-only access to a virtual machine, vRealize Operations Manager applies the PowerUser role to the user to access the virtual machine.

Refreshing Permissions

When you change permissions for a vCenter Server user in vCenter Server, the user must log out and log back in to vRealize Operations Manager to refresh the permissions and view the updated results in vRealize Operations Manager. Alternatively, the user can wait for vRealize Operations Manager to refresh. The permissions refresh at fixed intervals, as defined in the `$ALIVE_BASE/user/conf/auth.properties` file. The default refreshing interval is half an hour. If necessary, you can change this interval for all nodes in the cluster.

Single Sign-On and vCenter Users

When vCenter Server users log into vRealize Operations Manager by way of single sign-on, they are registered on the vRealize Operations Manager User Accounts page. If you delete the account of a vCenter Server user that has logged into vRealize Operations Manager by way of single sign-on, or remove the user from a single sign-on group, the user account entry still appears on the User Account page and you must delete it manually.

Generating Reports

vCenter Server users cannot create or schedule reports in vRealize Operations Manager.

Backward Compatibility for vCenter Server Users in vRealize Operations Manager

vRealize Operations Manager provides backward compatibility for users of the earlier version of vRealize Operations Manager, so that users of vCenter Server who have privileges in the earlier version in vCenter Server can log in to vRealize Operations Manager.

When you register vRealize Operations Manager in vCenter Server, certain roles become available in vCenter Server.

- The Administrator account in the previous version of vRealize Operations Manager maps to the PowerUser role.
- The Operator account in the previous version of vRealize Operations Manager maps to the ReadOnly role.

During registration, all roles in vRealize Operations Manager, except for vRealize Operations Manager Administrator, Maintenance, and Migration, become available dynamically in vCenter Server. Administrators in vCenter Server have all of the roles in vRealize Operations Manager that map during registration, but these administrator accounts only receive a specific role on the root folder in vCenter Server if it is specially assigned.

Registration of vRealize Operations Manager with vCenter Server is optional. If users choose not to register vRealize Operations Manager with vCenter Server, a vCenter Server administrator can still use their user name and password to log in to vRealize Operations Manager, but these users cannot use the vCenter Server session ID to log in. In this case, typical vCenter Server users must have one or more vRealize Operations Manager roles to log in to vRealize Operations Manager.

When multiple instances of vCenter Server are added to vRealize Operations Manager, user credentials become valid for all of the vCenter Server instances. When a user logs in to vRealize Operations Manager, if the user selects all vCenter Server options during login, vRealize Operations Manager requires that the user's credentials are valid for all of the vCenter Server instances. If a user account is only valid for a single vCenter Server instance, that user can select the vCenter Server instance from the login drop-down menu to log in to vRealize Operations Manager.

vCenter Server users who log in to vRealize Operations Manager must have one or more of the following roles in vCenter Server:

- vRealize Operations Content Admin Role
- vRealize Operations General User Role 1
- vRealize Operations General User Role 2
- vRealize Operations General User Role 3
- vRealize Operations General User Role 4
- vRealize Operations Power User Role
- vRealize Operations Power User without Remediation Actions Role
- vRealize Operations Read Only Role

For more information about vCenter Server users, groups, and roles, see the vCenter Server documentation.

External User Sources in vRealize Operations Manager

You can obtain user accounts from external sources so that you can use them in your vRealize Operations Manager instance.

There are two types of external user identity sources:

- Lightweight Directory Access Protocol (LDAP): Use the LDAP source if you want to use the Active Directory or LDAP servers as authentication sources. The LDAP source does not support multi-domains even when there is a two-way trust between Domain A and Domain B.
- Single Sign-On (SSO): Use a single sign-on source to perform single sign-on with any application that supports vCenter single sign-on, including vRealize Operations Manager. For example, you can install a standalone vCenter Platform Services Controller (PSC) and use it to communicate with an Active Directory server. Use a PSC if the Active Directory has a setup that is too complex for the simple LDAP source in vRealize Operations Manager, or if the LDAP source is experiencing slow performance.

Roles and Privileges in vRealize Operations Manager

vRealize Operations Manager provides several predefined roles to assign privileges to users. You can also create your own roles.

You must have privileges to access specific features in the vRealize Operations Manager user interface. The roles associated with your user account determine the features you can access and the actions you can perform.

Each predefined role includes a set of privileges for users to perform, create, read, update, or delete actions on components such as dashboards, reports, administration, capacity, policies, problems, symptoms, alerts, user account management, and adapters. For information about roles and associated permissions, see [KB 59484](#).

Administrator

Includes privileges to all features, objects, and actions in vRealize Operations Manager.

PowerUser

Users have privileges to perform the actions of the Administrator role except for privileges to user management and cluster management. vRealize Operations Manager maps vCenter Server users to this role.

PowerUserMinusRemediation

Users have privileges to perform the actions of the Administrator role except for privileges to user management, cluster management, and remediation actions.

ContentAdmin

Users can manage all content, including views, reports, dashboards, and custom groups in vRealize Operations Manager.

AgentManager

Users can deploy and configure End Point Operations Management agents.

GeneralUser-1 through GeneralUser-4

These predefined template roles are initially defined as ReadOnly roles. vCenter Server administrators can configure these roles to create combinations of roles to give users multiple types of privileges. Roles are synchronized to vCenter Server once during registration.

ReadOnly

Users have read-only access and can perform read operations, but cannot perform write actions such as create, update, or delete.

User Scenario: Manage User Access Control

As a system administrator or virtual infrastructure administrator, you manage user access control in vRealize Operations Manager so that you can ensure the security of your objects. Your company just hired a new person, and you must create a user account and assign a role to the account so that the new user has permission to access specific content and objects in vRealize Operations Manager.

In this scenario you will learn how to create user accounts and roles, and assign roles to the user accounts to specify access privileges to views and objects. You will then demonstrate the intended behavior of the permissions on these accounts.

You will create a new user account, named Tom User, and a new role that grants administrative access to objects in the vRealize Operations Clusters. You will apply the new role to the user account.

Finally, you will import a user account from an external LDAP user database that resides on another machine to vRealize Operations Manager, and assign a role to the imported user account to configure the user's privileges.

Prerequisites

Verify that the following conditions are met:

- vRealize Operations Manager is installed and operating properly, and contains objects such as clusters, hosts, and virtual machines.
- One or more user groups are defined.

What to do next

Create a new role.

Create a New Role

You use roles to manage access control for user accounts in vRealize Operations Manager.

In this procedure, you will add a new role and assign administrative permissions to the role.

Prerequisites

Verify that you understand the context of this scenario. See [User Scenario: Manage User Access Control](#). For information about roles and associated permissions, see [KB 59484](#).

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the **Roles** tab.
- 3 Click the **Add** icon on the toolbar to create a role.

The **Create Role** dialog box appears.

- 4 For the role name, type **admin_cluster**, then type a description and click **OK**.

The `admin_cluster` role appears in the list of roles.

- 5 Click the **admin_cluster** role.

- 6 In the Details grid below, on the Permissions pane, click the **Edit** icon.

The **Assign Permissions to Role** dialog box appears.

- 7 Select the **Administrative Access - all permissions** check box.

- 8 Click **Update**.

This action gives this role administrative access to all the features in the environment.

What to do next

Create a user account, and assign this role to the account.

Create a User Account

As an administrator you assign a unique user account to each user so that they can use vRealize Operations Manager. While you set up the user account, you assign the privileges that determine what activities the user can perform in the environment, and upon what objects.

In this procedure, you will create a user account, assign the `admin_cluster` role to the account, and associate the objects that the user can access while assigned this role. You will assign access to objects in the vRealize Operations Cluster. Then, you will test the user account to confirm that the user can access only the specified objects.

Prerequisites

Create a new role. See [Create a New Role](#).

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the **User Accounts** tab.
- 3 Click the **Add** icon to create a new user account, and provide the information for this account.

Option	Description
User Name	Type the user name to use to log in to vRealize Operations Manager.
Password	Type a password for the user.
Confirm Password	Type the password again to confirm it.
First Name	Type the user's first name. For this scenario, type Tom .
Last Name	Type the user's last name. For this scenario, type User .
Email Address	(Optional). Type the user's email address.
Description	(Optional). Type a description for this user.

Option	Description
Disable this user	Do not select this check box, because you want the user to be active for this scenario.
Require password change at next login	Do not select this check box, because you do not need to change the user's password for this scenario.

4 Click **Next**.

The list of user groups appears.

5 Select a user group to add the user account as a member of the group.

6 Click the **Objects** tab.

7 Select the **admin_cluster** role from the drop-down menu.

8 Select the **Assign this role to the user** check box.

9 In the Object Hierarchies list, select the **vRealize Operations Cluster** check box.

10 Click **Finish**.

You created a new user account for a user who can access all the vRealize Operations Cluster objects. The new user now appears in the list of user accounts.

11 Log out of vRealize Operations Manager.

12 Log in to vRealize Operations Manager as Tom User, and verify that this user account can access all the objects in the vRealize Operations Cluster hierarchy, but not other objects in the environment.

13 Log out of vRealize Operations Manager.

Results

You used a specific role to assign permission to access all objects in the vRealize Operations Cluster to a user account named Tom User.

What to do next

Import a user account from an external LDAP user database that resides on another machine, and assign permissions to the user account.

Import a User Account and Assign Permissions

You can import user accounts from external sources, such as an LDAP database on another machine, or a single sign-on server, so that you can give permission to those users to access certain features and objects in vRealize Operations Manager.

Prerequisites

- Configure an authorization source. See the vRealize Operations Manager Information Center.

Procedure

- 1 Log out of vRealize Operations Manager, then log in as a system administrator.
- 2 In the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 3 On the toolbar, click the **Import Users** icon.
- 4 Specify the options to import user accounts from an authorization source.
 - a On the Import Users page, from the **Import From** drop-down menu, select an authentication source.
 - b In the **Domain Name** drop-down menu, type the domain name from which you want to import users, and click **Search**.
 - c Select the users you want to import, and click **Next**.
 - d On the **Groups** tab, select the user group to which you want to add this user account.
 - e Click the **Objects** tab, select the **admin_cluster** role, and select the **Assign this role to the user** check box.
 - f In the Object Hierarchies list, select the **vRealize Operations Cluster** check box, and click **Finish**.
- 5 Log out of vRealize Operations Manager.
- 6 Log in to vRealize Operations Manager as the imported user.
- 7 Verify that the imported user can access only the objects in the vRealize Operations Cluster.

Results

You imported a user account from an external user database or server to vRealize Operations Manager, and assigned a role and the objects the user can access while holding this role to the user.

You have finished this scenario.

Configure a Single Sign-On Source in vRealize Operations Manager

As a system administrator or virtual infrastructure administrator, you use single sign-on to enable SSO users to log in securely to your vRealize Operations Manager environment.

After the single sign-on source is configured, users are redirected to an SSO identity source for authentication. When logged in, users can access other vSphere components such as the vCenter Server without having to log in again.

Prerequisites

- Verify that the server system time of the single sign-on source and vRealize Operations Manager are synchronized. If you need to configure the Network Time Protocol (NTP), see information about cluster and node maintenance in the *vRealize Operations Manager vApp Deployment and Configuration Guide*.

- Verify that you have access to a Platform Services Controller through the vCenter Server. See the VMware vSphere Information Center for more details.

Procedure

- 1 Log in to vRealize Operations Manager as an administrator.
- 2 In the menu, click **Administration**, then in the left pane click **Access > Authentication Sources**.
- 3 Click **Add**.
- 4 In the Add Source for User and Group Import dialog box, provide information for the single sign-on source.

Option	Action
Source Display Name	Type a name for the import source.
Source Type	Verify that SSO SAML is displayed.
Host	Enter the IP address or FQDN of the host machine where the single sign-on server resides. If you enter the FQDN of the host machine, verify that every non-remote collector node in the vRealize Operations Manager cluster can resolve the single sign-on host FQDN.
Port	Set the port to the single sign-on server listening port. By default, the port is set to 443.
User Name	Enter the user name that can log into the SSO server.
Password	Enter the password.
Grant administrator role to vRealize Operations Manager for future configuration?	Select Yes so that the SSO source is reregistered automatically if you make changes to the vRealize Operations Manager setup. If you select No , and the vRealize Operations Manager setup is changed, single sign-on users will not be able to log in until you manually reregister the single sign-on source.
Automatically redirect to vRealize Operations single sign-on URL?	Select Yes to direct users to the vCenter single-sign on log in page. If you select No , users are not redirected to SSO for authentication. This option can be changed in the vRealize Operations Manager Global Settings.
Import single sign-on user groups after adding the current source?	Select Yes so that the wizard directs you to the Import User Groups page when you have completed the SSO source setup. If you want to import user accounts, or user groups at a later stage, select No .
Advanced options	If your environment uses a load balancer, enter the IP address of the load balancer.

- 5 Click **Test** to test the source connection, and then click **OK**.
The certificate details are displayed.
- 6 Select the **Accept this Certificate** check box, and click **OK**.

- 7 In the Import User Groups dialog box, import user accounts from an SSO server on another machine.

Option	Action
Import From	Select the single sign-on server you specified when you configured the single sign-on source.
Domain Name	Select the domain name from which you want to import user groups. If Active Directory is configured as the LDAP source in the PSC, you can only import universal groups and domain local groups if the vCenter Server resides in the same domain.
Result Limit	Enter the number of results that are displayed when the search is conducted.
Search Prefix	Enter a prefix to use when searching for user groups.

- 8 In the list of user groups displayed, select at least one user group, and click **Next**.
- 9 In the Roles and Objects pane, select a role from the **Select Role** drop-down menu, and select the **Assign this role to the group** check box.
- 10 Select the objects users of the group can access when holding this role.
To assign permissions so that users can access all the objects in vRealize Operations Manager, select the **Allow access to all objects in the system** check box.
- 11 Click **OK**.
- 12 Familiarize yourself with single-sign on and confirm that you have configured the single sign-on source correctly.
 - a Log out of vRealize Operations Manager.
 - b Log in to the vSphere Web Client as one of the users in the user group you imported from the single sign-on server.
 - c In a new browser tab, enter the IP address of your vRealize Operations Manager environment.
 - d If the single sign-on server is configured correctly, you are logged in to vRealize Operations Manager without having to enter your user credentials.

Edit a Single Sign-On Source

Edit a single sign-on source if you need to change the administrator credentials used to manage the single sign-on source, or if you have changed the host of the source.

When you configure an SSO source, you specify either the IP address or the FQDN of the host machine where the single sign-on server resides. If you want to configure a new host, that is, if the single sign-on server resides on a different host machine than the one configured when the source was set up, vRealize Operations Manager removes the current SSO source, and creates a new source. In this case, you must reimport the users you want to associate with the new SSO source.

If you want to change the way the current host is identified in vRealize Operations Manager, for example, change the IP address to the FQDN and the reverse, or update the IP address of the PSC if the IP address of the configured PSC has changed, vRealize Operations Manager updates the current SSO source, and you are not required to reimport users.

Procedure

- 1 Log in to vRealize Operations Manager as an administrator.
- 2 In the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.
- 3 Select the single sign-on source and click the **Edit** icon.
- 4 Make changes to the single sign-on source, and click **OK**.

If you are configuring a new host, the New Single Sign-On Source Detected dialog box appears.

- 5 Enter the administrator credentials that were used to set up the single sign-on source, and click **OK**.

The current SSO source is removed, and a new one created.

- 6 Click **OK** to accept the certificate.
- 7 Import the users you want to associate with the SSO source.

Audit Users and the Environment in vRealize Operations Manager

At times you might need to provide documentation as evidence of the sequence of activities that took place in your vRealize Operations Manager environment. Auditing allows you to view the users, objects, and information that is collected. To meet audit requirements, such as for business critical applications that contain sensitive data that must be protected, you can generate reports on the activities of your users, the privileges assigned to users to access objects, and the counts of objects and applications in your environment.

Auditing reports provide traceability of the objects and users in your environment.

User Activity Audit

Run this report to understand the scope of user activities, such as logging in, actions on clusters and nodes, changes to system passwords, activating certificates, and logging out.

User Permissions Audit

Generate this report to understand the scope of user accounts and their roles, access groups, and access privileges.

System Audit

Run this report to understand the scale of your environment. This report displays the counts of configured and collecting objects, the types and counts of adapters, configured and collecting metrics, super metrics, applications, and existing virtual environment objects. This

report can help you determine whether the number of objects in your environment exceeds a supported limit.

System Component Audit

Run this report to display a version list of all the components in your environment.

Reasons for Auditing Your Environment

Auditing in vRealize Operations Manager helps data center administrators in the following types of situations.

- You must track each configuration change to an authenticated user who initiated the change or scheduled the job that performed the change. For example, after an adapter changes an object, which is associated with a specific object identifier at a specific time, the data center administrator can determine the principal identifier of the authenticated user who initiated the change.
- You must track who made changes to your data center during a specific range of time, to determine who changed what on a particular day. You can identify the principal identifiers of authenticated users who were logged in to vRealize Operations Manager and running jobs, and determine who initiated the change.
- You must determine which objects were affected by a particular user during a time specific range of time.
- You must correlate events that occurred in your data center, and view these events overlaid so that you can visualize relationships and the cause of the events. Events can include login attempts, system startup and shutdown, application failures, watchdog restarts, configuration changes of applications, changes to security policy, requests, responses, and status of success.
- You must validate that the components installed in your environment are running the latest version.

System Component Audit

A system component audit report provides a version list of every component installed in the system.

Where You Audit System Components

- 1 To audit system components, in the menu, click **Administration**, and then in the left pane click **History > Audit**.
- 2 Click the **System Component Audit** tab.

A list of components installed in the environment appears on the page.

Table 9-1. System Component Audit Actions

Option	Description
Download	Display the version information in a new browser window.

vRealize Operations Manager Passwords and Certificates

For secure vRealize Operations Manager operation, you might need to perform maintenance on passwords or authentication certificates.

- Passwords are for user access to the product interfaces or to console sessions on cluster nodes.
- Authentication certificates are for secure machine-to-machine communication within vRealize Operations Manager itself or between vRealize Operations Manager and other systems.

Reset the vRealize Operations Manager Administrator Password

You might need to reset the vRealize Operations Manager administrator password as part of securing or maintaining your deployment and if you forget the admin account password.

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at `https://<master-node-name> or <master-node-ip-address>/admin`.
- 2 Log in with the admin user name and password for the master node.
- 3 In the left pane, click **Administrator Settings**.
- 4 In the **Change Administrator Password** section, enter the current password, and enter the new password twice to ensure its accuracy.

Note You cannot change the administrator user name.

- 5 Click **Save**.
- 6 Optionally, to recover a forgotten password, configure the **Password Recovery Settings**.

Table 9-2. Password Recovery Settings

Password Recovery Settings Options	Description
Your E-mail	Email id to which you want to receive the recovery email.
SMTP Server	smtp.vmware.com
Port	Port used for the communication. By default, 25 is used for a non secure port and 465 for a secure port.
SSL (SMTPS)	Enable or disable to protect the communication using the secure socket layer.
STARTTLS Encryption	Enable or disable to switch the insecure communication starting with the TLS handshake.

Table 9-2. Password Recovery Settings (continued)

Password Recovery Settings Options	Description
Sender E-mail	The email from which the password recovery email is sent.
User name	Username for the SMTP server account, as some servers require authentication.
Password	Password for the SMTP server account.
Test	To verify the mandatory fields and make an attempt to communicate with the given SMTP server.

- 7 Click **Save**. Optionally, click **Reset** to enter the details again.

Generate a vRealize Operations Manager Passphrase

When users need to add a node to the vRealize Operations Manager cluster, you can generate a temporary passphrase instead of giving them the primary administrator login credentials, which might be a security issue.

A temporary passphrase is good for one use only.

Prerequisites

Create and configure the primary node.

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at <https://master-node-name-or-ip-address/admin>.
- 2 Log in with the admin user name and password for the master node.
- 3 In the list of cluster nodes, select the master node.
- 4 From the toolbar above the list, click the option to generate a passphrase.
- 5 Enter a number of hours before the passphrase expires.
- 6 Click **Generate**.

A random alphanumeric string appears, which you can send to a user who needs to add a node.

What to do next

Have the user supply the passphrase when adding a node.

Custom vRealize Operations Manager Certificates

By default, vRealize Operations Manager includes its own authentication certificates. The default certificates cause the browser to display a warning when you connect to the vRealize Operations Manager user interface.

Your site security policies might require that you use another certificate, or you might want to avoid the warnings caused by the default certificates. In either case, vRealize Operations Manager supports the use of your own custom certificate. You can upload your custom certificate during initial primary node configuration or later.

Custom vRealize Operations Manager Certificate Requirements

A certificate used with vRealize Operations Manager must conform to certain requirements. Using a custom certificate is optional and does not affect vRealize Operations Manager features.

Requirements for Custom Certificates

Custom vRealize Operations Manager certificates must meet the following requirements.

- The certificate file must include the terminal (leaf) server certificate, a private key, and all issuing certificates if the certificate is signed by a chain of other certificates.
- In the file, the leaf certificate must be first in the order of certificates. After the leaf certificate, the order does not matter.
- In the file, all certificates and the private key must be in PEM format. vRealize Operations Manager does not support certificates in PFX, PKCS12, PKCS7, or other formats.
- In the file, all certificates and the private key must be PEM-encoded. vRealize Operations Manager does not support DER-encoded certificates or private keys.

PEM-encoding is base-64 ASCII and contains legible BEGIN and END markers, while DER is a binary format. Also, file extension might not match encoding. For example, a generic `.cer` extension might be used with PEM or DER. To verify encoding format, examine a certificate file using a text editor.

- The file extension must be `.pem`.
- The private key must be generated by the RSA or DSA algorithm.
- The private key must not be encrypted by a pass phrase if you use the primary node configuration wizard or the administration interface to upload the certificate.
- The REST API in this vRealize Operations Manager release supports private keys that are encrypted by a pass phrase. Contact VMware Technical Support for details.
- The vRealize Operations Manager Web server on all nodes will have the same certificate file, so it must be valid for all nodes. One way to make the certificate valid for multiple addresses is with multiple Subject Alternative Name (SAN) entries.
- SHA1 certificates creates browser compatibility issues. Therefore, ensure that all certificates that are created and being uploaded to vRealize Operations Manager are signed using SHA2 or newer.
- The vRealize Operations Manager supports custom security certificates with key length up to 8192 bits. An error is displayed when you try to upload a security certificate generated with a stronger key length beyond 8192 bits.

For more information, see the following KB articles:

- [vRealize Operations Manager 6.x fails to accept and apply Custom CA Certificate \(2144949\)](#)

Configure a Custom Certificate

You can use OpenSSL to configure an authentication certificate for use with vRealize Operations Manager. You must first generate a Certificate PEM for vRealize Operations Manager, then install the Certificate PEM in vRealize Operations Manager. The certificates applied through the vRealize Operations Manager Admin UI will be used only for securely connecting and serving the user interfaces to (external) clients. We do not update the certificates for specific components of vRealize Operations Manager.

Procedure

1 Generate a Certificate PEM file for use with vRealize Operations Manager

- Generate a key pair by running this command:

```
openssl genrsa -out key_filename.key 2048
```

- Use the key to generate a certificate signing request by running this command:

```
openssl req -new -key key_filename.key -out certificate_request.csr
```

- Submit the CSR file to your Certificate Authority (CA) to obtain a signed certificate.
- From your Certificate Authority, download the certificate and the complete issuing chain (one or more certificates). Download them in Base64 format.
- Enter the command to create a single PEM file containing all certificates and the private key. In this step, the example certificate is *server_cert.cer* and the issuing chain is *cacerts.cer*.

Note The order of CA's certs in the .PEM file: Cert, Private Key, Intermediate Cert and then Root Cert.

```
cat server_cert.cer key_filename.key cacerts.cer > multi_part.pem
```

In Windows replace cat with type.

The finished PEM file should look similar to the following example, where the number of CERTIFICATE sections depends on the length of the issuing chain:

```
-----BEGIN CERTIFICATE-----
(Your Primary SSL certificate: your_domain_name.crt)
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
(Your Private Key: your_domain_name.key)
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
(Your Intermediate certificate: DigiCertCA.crt)
```



```
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
(Your Root certificate: TrustedRoot.crt)
-----END CERTIFICATE-----
```

2 Install a PEM in vRealize Operations Manager

- a In a Web browser, navigate to the vRealize Operations Manager administration interface.

```
https://vrops-node-FQDN-or-ip-address/admin
```

- b Log in with the admin username and password.
- c At the upper right, click the yellow **SSL Certificate** icon.
- d In the **SSL Certificate** window, click **Install New Certificate**.
- e Click **Browse** for certificate.
- f Locate the certificate .pem file, and click Open to load the file in the **Certificate Information** text box. The certificate file must contain a valid private key and a valid certificate chain.
- g Click **Install**.

Verifying a Custom vRealize Operations Manager Certificate

When you upload a custom certificate file, the vRealize Operations Manager interface displays summary information for all certificates in the file.

For a valid custom certificate file, you should be able to match issuer to subject, issuer to subject, back to a self-signed certificate where the issuer and subject are the same.

In the following example, OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32 is issued by OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32, which is issued by OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84, which is issued by itself.

```
Thumbprint: 80:C4:84:B9:11:5B:9F:70:9F:54:99:9E:71:46:69:D3:67:31:2B:9C
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32
Subject Alternate Name:
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:24.000Z
Valid To: 2020-05-06T16:25:24.000Z

Thumbprint: 72:FE:95:F2:90:7C:86:24:D9:4E:12:EC:FB:10:38:7A:DA:EC:00:3A
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:19.000Z
Valid To: 2020-05-06T16:25:19.000Z

Thumbprint: FA:AD:FD:91:AD:E4:F1:00:EC:4A:D4:73:81:DB:B2:D1:20:35:DB:F2
```

```

Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-
e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-
e1806f0d3f84
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:24:45.000Z
Valid To: 2020-05-06T16:24:45.000Z

```

Sample Contents of Custom vRealize Operations Manager Certificates

For troubleshooting purposes, you can open a custom certificate file in a text editor and inspect its contents.

PEM Format Certificate Files

A typical PEM format certificate file resembles the following sample.

```

-----BEGIN CERTIFICATE-----
MIIF1DCCBlygAwIBAgIKFYXYUwAAAAAGTANBgkqhkiG9w0BAQ0FADBhMRMwEQYK
CZImiZPyLQG8GRYDY29tMRUwEwYKCCImiZPyLQG8GRYFdm13Y3MxGDAWBgoJkiaJ
<snip>
vKStQJNr7z2+pTy92M6FgJz3y+daL+9ddbaMnp9fVXjHB0DLGGaL0vyD+KJ8+xba
aGJfGf9ELXM=
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQE4l5ffX694riI1RmdRLJwL6sOWa+Wf70HRoLtx21kZzbXbUQN
mQhTRiidJ3Ro2gRbj/btSsI+OMUotz5VVRT/yeyoTC5l2uJEapld45RroUDHQwWJ
<snip>
DAN9hQus3832xMkAuVP/jt76dHDYyviyIYbmxxMa1X7LZy1MCQVg4hCH0vLsHtLh
M1rOAsz62Eht/ib61AsVCCiN3gLRX7MKsYdxZcRVruGXSIh33ynA
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIDnTCCAowGawIBAgIQY+j29InmdYNCs2cK1H4kPzANBgkqhkiG9w0BAQ0FADBh
MRMwEQYKCCImiZPyLQG8GRYDY29tMRUwEwYKCCImiZPyLQG8GRYFdm13Y3MxGDAW
<snip>
ukzUuqX7wEhc+QgJWgl41mWZBZ09gfsA9XuXBL0k17IpVHpEgwwrjQz8X68m4I99
dD5Pf1f/nLRJvR9jwXl62yk=
-----END CERTIFICATE-----

```

Private Keys

Private keys can appear in different formats but are enclosed with clear BEGIN and END markers.

Valid PEM sections begin with one of the following markers.

```

-----BEGIN RSA PRIVATE KEY-----
-----BEGIN PRIVATE KEY-----

```

Encrypted private keys begin with the following marker.

```

-----BEGIN ENCRYPTED PRIVATE KEY-----

```

Bag Attributes

Microsoft certificate tools sometimes add Bag Attributes sections to certificate files. vRealize Operations Manager safely ignores content outside of BEGIN and END markers, including Bag Attributes sections.

```
Bag Attributes
Microsoft Local Key set: <No Values>
localKeyID: 01 00 00 00
Microsoft CSP Name: Microsoft RSA SChannel Cryptographic Provider
friendlyName: le-WebServer-8dea65d4-c331-40f4-aa0b-205c3c323f62
Key Attributes
X509v3 Key Usage: 10
-----BEGIN PRIVATE KEY-----
MIICdwIBADANBgkqhkiG9w0BAQEFAASCAmEwggJdAgEAAoGBAKHqyfc+qcQK4yxJ
om3PuB8dYzm34Qlt81GAAnBPYe3B4Q/0ba6PV8GtWG2svIpc1/eflwGHgTU3zJxR
gkKh7I3K5tGESn81ipyKtKbYebh+aBMqPKrNNUEKlr0M9sa3WSc0o3350tCc1ew
5ZkNYZ4BRUVYwm0HogeGh0thRn2fAgMBAAECGyABhPmGN3FSZKPDG6HJLARvTLBH
KAGVnBGHd0M0mMabghFBnBKXa8LwD1dgGBng1oOakEXTftkIjdB+uwkU5P4aRr07
vGuJUtRyRCU/4fjLBDuxQL/KpQfruAQaoF9uWUwh5W9fEeW3g26fzVL8AFZnbXS0
7Z0AL1H3LncLd5rpQJBANnI7vFu06bFxVF+kq6Z0JFMx7x3K4VGxgg+PfFEBEPS
UJ2LuDH5/Rc63BaxFzM/q3B3Jhehvgw61mMyxU7QSSUCQC+VDuW3XEWJjSiU6KD
gEGpCyJ5SBepBLsukljpgidKkDNlKlgbWVytCVkTAmuoAz33kMwfqIiNcqQbUgVV
UnpzAkB7d0CP00deSsy8kMdTmKXlkf4qSF0x55epYK/5MZhBYuA1ENrR6mmjW8ke
TDNc6IGm9sVvrFBz2n9kKYpWThrJAKeAK5R69DtW0cbkLy5MqEzOHQauP36gDi1L
WMXPvUfzSYTQ5aM2rrY2/1FtSSkqUwFYh9sw8eDbqVpIV4rc6dDfcwJBALiDPT0
tz86wySJNe0iUkQm36iXVF8AckPKT9TrbC3Ho7nC80zL7gEl1ETa4Zc86Z3wpcGF
BHhEDMHaihyuVgI=
-----END PRIVATE KEY-----

Bag Attributes
localKeyID: 01 00 00 00
1.3.6.1.4.1.311.17.3.92: 00 04 00 00
1.3.6.1.4.1.311.17.3.20: 7F 95 38 07 CB 0C 99 DD 41 23 26 15 8B E8
D8 4B 0A C8 7D 93
friendlyName: cos-oc-vcops
1.3.6.1.4.1.311.17.3.71: 43 00 4F 00 53 00 2D 00 4F 00 43 00 2D 00
56 00 43 00 4D 00 35 00 37 00 31 00 2E 00 76 00 6D 00 77 00 61 00
72 00 65 00 2E 00 63 00 6F 00 6D 00 00 00
1.3.6.1.4.1.311.17.3.87: 00 00 00 00 00 00 00 00 02 00 00 00 20 00
00 00 02 00 00 00 6C 00 64 00 61 00 70 00 3A 00 00 00 7B 00 41 00
45 00 35 00 44 00 44 00 33 00 44 00 30 00 2D 00 36 00 45 00 37 00
30 00 2D 00 34 00 42 00 44 00 42 00 2D 00 39 00 43 00 34 00 31 00
2D 00 31 00 43 00 34 00 41 00 38 00 44 00 43 00 42 00 30 00 38 00
42 00 46 00 7D 00 00 00 70 00 61 00 2D 00 61 00 64 00 63 00 33 00
2E 00 76 00 6D 00 77 00 61 00 72 00 65 00 2E 00 63 00 6F 00 6D 00
5C 00 56 00 4D 00 77 00 61 00 72 00 65 00 20 00 43 00 41 00 00 00
31 00 32 00 33 00 33 00 30 00 00 00
subject=/CN=cos-oc-vcops.eng.vmware.com
issuer=/DC=com/DC=vmware/CN=VMware CA
-----BEGIN CERTIFICATE-----
MIIFWTCCBEGgAwIBAgIKSjGT5gACAAAwKjANBgkqhkiG9w0BAQUFADBMMwEQYK
CZImiZPYLQBGRYDY29tMRYwFAYKCZImiZPYLQBGRYGdm13YXJlMRIwEAYDVQQD
EwltWlTXdhcmUgQ0EwHhcNMjQwMjA1MTg1OTM2WhcNMjYwMjA1MTg1OTM2WjAmMSQw
```

Modifying Global Settings

The global settings control the system settings for vRealize Operations Manager, including data retention and system timeout settings. You can modify one or more of the settings to monitor your environment better. These settings affect all your users.

The global settings do not affect metric interactions, color indicators, or other object management behaviors. These behaviors are configured in your policies.

Settings related to managing objects with vRealize Operations Manager are available on the **Inventory Explorer** page.

You can view tooltips for each option in the Edit Global Settings dialog box.

Global Settings Best Practices

Most of the settings pertain to how long vRealize Operations Manager retains collected and process data.

The default values are common retention periods. You might need to adjust the time periods based on your local policies or disk space.

List of Global Settings

The global settings determine how vRealize Operations Manager retains data, keeps connection sessions open, and other settings. These are system settings that affect all users.

Table 9-3. Global Setting Default Values and Descriptions

Setting	Default Value	Description
Action History	30 days	Number of days to retain the recent task data for actions. The data is purged from the system after the specified number of days.
Deleted Objects	168 hours	Number of hours to retain objects that are deleted from an adapter data source or server before deleting them from vRealize Operations Manager. An object deleted from an adapter data source is identified by vRealize Operations Manager as not existing and vRealize Operations Manager can no longer collect data about the object. Whether vRealize Operations Manager identifies deleted objects as not existing depends the adapter. This feature is not implemented in some adapters. For example, if the retention time is 360 hour and a virtual machine is deleted from a vCenter Server instance, the virtual machine remains as an object in vRealize Operations Manager for 15 days before it is deleted. This setting applies to objects deleted from the data source or server, not to any objects you delete from vRealize Operations Manager on the Inventory Explorer page. A value of -1 deletes objects immediately.

Table 9-3. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Deletion Schedule Interval	24 hours	Determines the frequency to schedule deletion of resources. This setting works with the Deleted Objects setting to remove objects that no longer exist in the environment. vRealize Operations Manager transparently marks objects for removal that have not existed for the length of time specified under Deleted Objects. vRealize Operations Manager then removes the marked objects at the frequency specified under Deletion Scheduling Interval.
Object History	90 days	Number of days to retain the history of the object configuration, relationship, and property data. The configuration data is the collected data from the monitored objects on which the metrics are based. The collected data includes changes to the configuration of the object. The data is purged from the system after the specified number of days.
Session Timeout	30 minutes	If your connection to vRealize Operations Manager is idle for the specified amount of time, you are logged out of the application. You must provide credentials to log back in.
Symptoms/Alerts	45 days	Number of days to retain canceled alerts and symptoms. The alerts and symptoms are either canceled by the system or by a user.
Time Series Data Retention	6 months	Number of months that you want to retain the collected and calculated metric data for the monitored objects. This setting is set to 6 months by default for 5 minutes interval data retention.
Additional Time Series Data Retention	36 months	The number of months that the roll-up data extends beyond the regular period. The roll-up data is available starting from the end of the regular period and until the end of the roll-up data retention period. If you specify 0 as the value, then this will effectively disable the Additional Time Series Data Retention time and only data specified in Time Series Retention is stored. This setting ensures that after 6 months of normal retention for 5 minutes, the seventh month data is rolled up into a one Hour roll up. You can set up this option up to 120 months for data roll ups.
Deleted Users	100 days	You can specify the number of days to keep custom content created by a user who has been removed from vRealize Operations Manager or by the automatic synchronization of LDAP. For example, the custom dashboards created by a user.
Maintain Relationship History		You can maintain a history of all the relationships of all the monitored objects in vRealize Operations Manager.

Table 9-3. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Dynamic Threshold Calculation	enabled	<p>Determines whether to calculate normal levels of threshold violation for all objects.</p> <p>If the setting is disabled, the following area of vRealize Operations Manager does not work or are not displayed:</p> <ul style="list-style-type: none"> ■ Alert symptom definitions based on dynamic thresholds will not work ■ Metric charts that display normal behavior are not present <p>Disable this setting only if you have no alternative options for managing resource constraints for your vRealize Operations Manager system.</p>
Cost Calculation		The host time at which cost calculations are run.
Allow vCenter users to log in to individual vCenters using the vRealize Operations Manager UI		<p>Determine how users of vCenter Server login to vRealize Operations Manager.</p> <ul style="list-style-type: none"> ■ In the vRealize Operations Manager user interface, vCenter Server users can log in to individual vCenter Server instances. Disabled by default. ■ vCenter Server users can log in from vCenter Server clients. Enabled by default. ■ In the vRealize Operations Manager user interface, vCenter Server users can log in to all vCenter Server instances. Enabled by default.
Allow vCenter users to log in from vCenter clients		
Allow vCenter users to log in to all vCenters using the vRealize Operations Manager UI		
Automated Actions	enabled or disabled	Determines whether to allow vRealize Operations Manager to automate actions. When an alert triggered, the alert provides recommendations for remediation. You can automate an action to remediate an alert when the recommendation is the first priority for that alert. You enable actionable alerts in your policies.

Table 9-3. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Enable Standard Certification Validation		<p>This option enables certificate verification to Test Connection in the Create or Modify AI screen, using a standard verification flow.</p> <p>The option checks CA authority.</p> <ul style="list-style-type: none"> ■ Certificate Subject DN ■ Subject alternative name ■ Certificate validity period ■ Revocation list <p>This option also presents dialogs to user if one of those checks fail. It is up to the adapter implementation on how the adapter checks source certificate validity during a normal collection cycle. On a usual scenario, adapters just perform a thumb-print verification. However, in case this flag is enabled, Test connection validates certificates in full scale and accepts certificates that are matching all criteria without any user dialogs.</p>
Currency		<p>You can specify the currency unit that is used for all the cost calculations. You can select the type of currency from the list of currency types by clicking Choose Currency. From the Set Currency, select the required currency and confirm your action by clicking the check box, and set the currency.</p>
Customer Experience Improvement Program	enabled	<p>Determines whether to participate in the Customer Experience Improvement Program by having vRealize Operations Manager send anonymous usage data to https://vmware.com.</p>

Global Settings

To manage how vRealize Operations Manager retains data, keeps connection sessions open, and other settings, you can modify the values for the global settings. These system settings affect all users.

You can also choose to participate in the customer experience improvement program. For more information on accessing Global settings, see [Access Global Settings](#).

Access Global Settings

With global settings, you set times to delete objects, set timeouts, store historical data, use dynamic threshold and capacity calculations, and determine how vCenter Server users log in. For automated actions, you can select whether to allow actions to be triggered from alert recommendations automatically.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Management > Global Settings**.

- 2 To edit the global settings, click the Edit icon.

Table 9-4. Global Settings Options

Option	Description
Edit Global Settings	Use the toolbar option to modify setting values.
Setting	Setting name.
Value	Current value for the setting. To change the setting value, click Edit Global Settings .
Description	Information about the setting. Point to the setting to display additional information about the setting.

Transfer Ownership of Dashboards and Report Schedules

When a user is deleted from vRealize Operations Manager, the report schedules and dashboards created by the user are stored as orphaned content. As an admin user, you can transfer ownership of dashboards and report schedules created by deleted users.

From Where You Can Transfer Ownership of Dashboards and Report Schedules

In the menu, click **Administration**. From the left pane, select **Management > Orphaned Content**.

Orphaned Content Page

You can view a list of deleted users from the **Deleted Users** panel in the left pane of the **Orphaned Content** page. Based on your selection in the **Deleted Users** panel, the dashboards and report schedules for the deleted user are displayed under the **Dashboard** and **Report Schedules** tabs in the **Orphaned Content** page.

As an admin user, you can take ownership, assign ownership, or discard orphaned dashboards and report schedules, from the **Actions** menu in the **Dashboards** and **Report Schedules** tabs. Enter the name or part of the name of a dashboard or report schedule in the **Filter** option and click **Enter**. The relevant dashboard or report schedule is displayed.

Table 9-5. Actions Menu Options

Actions	Options
Take Ownership	You can take ownership of the selected dashboards or report schedules.
Assign Ownership	You can assign a new owner for the selected dashboards or report schedules. You can select a target user from the Transfer Dashboards/Report Schedule dialog box.
Discard	You can permanently delete the dashboards or report schedules.

Create a vRealize Operations Manager Support Bundle

You create a vRealize Operations Manager support bundle to gather log and configuration files for analysis when troubleshooting a vRealize Operations Manager issue.

When you create a support bundle, vRealize Operations Manager gathers files from cluster nodes into ZIP files for convenience.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Support > Support Bundles**.
- 2 From the toolbar, click the **Create a Support Bundle** icon.
- 3 Select the option to create a **Light** or **Full support bundle**.
- 4 Select the cluster nodes that need to be evaluated for support.

Only logs from the selected nodes are included in the support bundle.

- 5 Click **OK**, and click **OK** to confirm support bundle creation.

Depending on the size of the logs and number of nodes, it might take time for vRealize Operations Manager to create the support bundle.

What to do next

Use the toolbar to download the support bundle ZIP files for analysis. For security, vRealize Operations Manager prompts you for credentials when you download a support bundle.

You can review the log files for error messages or, if you need troubleshooting assistance, send the diagnostic data to VMware Technical Support. When you resolve or close the issue, use the toolbar to delete the outdated support bundle to save disk space.

Customizing Icons

Every object or adapter in your environment has an icon representation. You can customize how the icon appears.

vRealize Operations Manager assigns a default icon to each object type and adapter type. Taken collectively, object types and adapter types are known as objects in your environment. Icons represent objects in the UI and help you to identify the type of object. For example, in the Topology Graph widget on a dashboard, labeled icons show how objects are connected to one other. You can quickly identify the type of object from the icon.

If you want to differentiate objects, you can change the icon. For example, a virtual machine icon is generic. If you want to pictorially distinguish the data that a vSphere virtual machine provides from the data that a Hypervisor virtual machine provides, you can assign a different icon to each.

Customize an Object Type Icon

You can use the default icons that vRealize Operations Manager provides, or you can upload your own graphics file for an object type. When you change an icon, your changes take effect for all users.

Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons**.
- 2 Click the **Object Type Icons** tab.
- 3 Assign the Object Type icon.
 - a Select the object type in the list with the icon to change.

By default, object types for all adapter types are listed. To limit the selection to the object types that are valid for a single adapter type, select the adapter type from the drop-down menu.
 - b Click the **Upload** icon.
 - c Browse to and select the file to use and click **Done**.
- 4 (Optional) To return to the default icon, select the object type and click the **Assign Default Icons** icon.

The original default icon appears.

Customize an Adapter Type Icon

You can use the default icons that vRealize Operations Manager provides, or you can upload your own graphics file for an adapter type. When you change an icon, your changes take effect for all users.

Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons**.
- 2 Click **Adapter Type Icons** tab.

- 3 Assign the Adapter Type icon.
 - a Select the adapter type in the list with the icon to change.
 - b Click the **Upload** icon.
 - c Browse to and select the file to use and click **Done**.
- 4 (Optional) To return to the default icon, select the adapter type and click the **Assign Default Icons** icon.

The original default icon appears.

OPS-CLI Command-Line Tool

10

The OPS-CLI tool is a Java application that you can use to manipulate the vRealize Operations Manager database. It replaces the VCOPS-CLI and DBCLI tools.

The product includes the executable file in the tools directory or in `<VCOPS_BASE>/tools/opscli/`.

Operating System	Filename
Linux	ops-cli.sh
Python	ops-cli.py

All OPS-CLI commands use the `-h` parameter for interactive and localized help.

When you add the `control` command to the `post_install.sh` script, it triggers the `redescribe` process after an adapter is installed or upgraded.

```
control -h | redescribe --force
```

Related Command-Line Documentation

In addition to the OPS-CLI, the VMware PowerCLI provides an easy-to-use Windows PowerShell interface for command-line access to administration tasks or for creating executable scripts. Documentation for all VMware PowerCLI versions is available from [VMware PowerCLI](#).

Supported Operations

The OPS-CLI tool supports the following database operations.

- [dashboard Command Operations](#)

You use the `dashboard` command to import, export, share, unshare, delete, reorder, show, hide, and set the default summary for dashboards.

- [template Command Operations](#)

You use the `template` command to import, export, share, unshare, delete, and reorder templates.

■ [supermetric Command Operations](#)

You use the `supermetric` command to import, export, configure, and delete super metrics.

■ [attribute Command Operations](#)

You use the `attribute` command to configure properties of a specific metric in one or more packages. The metric is the object attribute.

■ [reskind Command Operations for Object Types](#)

You use the `reskind` command to configure the default settings in your object type as defined by the `ResourceKind` model element. The command sets the default attribute or supermetric package, enables or disables dynamic thresholds, and enables or disables early warning smart alerts.

■ [report Command Operations](#)

You use the `report` command to import, export, configure, and delete super metrics.

■ [view Command Operations](#)

You use the `view` command to import, export, or delete view definitions.

■ [file Command Operations](#)

You use the `file` command to import, export, list, or delete database files. The command operates on metric, text widget, and topology widget files.

dashboard Command Operations

You use the `dashboard` command to import, export, share, unshare, delete, reorder, show, hide, and set the default summary for dashboards.

The `dashboard` command uses the following syntax.

```
dashboard -h | import|defsummary|export|share|unshare|delete|reorder|show|hide [parameters]
```

Table 10-1. dashboard Command Options

Command Name	Description	Syntax
dashboard import	Import a dashboard from a file and assign the ownership to a user account.	<pre>dashboard import -h user-name all group:group_name input- file [--force] [--share all group-name[{,group- name}]] [--retry maxRetryMinutes] [--set rank] [--default] [--create]</pre>
dashboard export	Export an existing dashboard to a file.	<pre>dashboard export -h user-name dashboard-name [output-dir]</pre>
dashboard defsummary	Import a dashboard from a file and assign the ownership to a user account.	<pre>dashboard defsummary -h input-file default --adapterKind adapterKind -- resourceKind resourceKind</pre>

Table 10-1. dashboard Command Options (continued)

Command Name	Description	Syntax
dashboard share	Share an existing dashboard with one or multiple user groups.	<code>dashboard share -h user-name dashboard-name all group-name[,{group-name}]</code>
dashboard unshare	Stop sharing a dashboard with specified groups.	<code>dashboard unshare -h user-name dashboard-name all group-name[,{group-name}]</code>
dashboard delete	Permanently delete a dashboard.	<code>dashboard delete -h user-name all group:group_name dashboard-name</code>
dashboard reorder	Set the order rank for a dashboard, with an option to make it the default.	<code>dashboard reorder -h user-name all group:group_name dashboard-name [--set rank] [--default]</code>
dashboard show	Show a dashboard.	<code>dashboard show -h user-name all group:group_name {,dashbaordname} all</code>
dashboard hide	Hide a dashboard.	<code>dashboard hide -h user-name all group:group_name {,dashboardname} all</code>

template Command Operations

You use the `template` command to import, export, share, unshare, delete, and reorder templates.

The `template` command uses the following syntax.

```
template -h | import|export|share|unshare|delete|reorder [parameters]
```

Table 10-2. template Command Operations

Command Name	Description	Syntax
template import	Import a template from a file.	<code>template import -h input-file [--force] [--share all group-name[,{group-name}]] [--retry maxRetryMinutes] [--set rank] [--create]</code>
template export	Export an existing template to a template file.	<code>template export -h template-name [output-dir]</code>
template share	Share an existing template with one or multiple user groups.	<code>template share -h template-name all group-name[,{group-name}]</code>

Table 10-2. template Command Operations (continued)

Command Name	Description	Syntax
template unshare	Stop sharing a template with specified groups.	<code>template unshare -h template-name all group-name[{{,group-name}}]</code>
template delete	Permanently delete a template.	<code>template delete -h template-name</code>
template reorder	Set the order rank for a template. The order rank controls the order of templates created based on shared templates.	<code>template reorder -h template-name [--set rank]</code>

supermetric Command Operations

You use the supermetric command to import, export, configure, and delete super metrics.

The supermetric command uses the following syntax.

```
supermetric -h | import|export|configure|delete [parameters]
```

Table 10-3. supermetric Command Operations

Command Name	Description	Syntax
supermetric import	Import a super metric from a file and assign the ownership to the specific user account.	<code>supermetric import -h input-file [--force] [--policies all policy-name[{{,policy-name}}] [--check (true false)] [--retry maxRetryMinutes] [--create]</code>
supermetric export	Export an existing super metric to a template file.	<code>supermetric export -h supermetric-name [output-dir]</code>

Table 10-3. supermetric Command Operations (continued)

Command Name	Description	Syntax
supermetric configures	Configure properties of a super metric in one or more super metrics packages.	<pre>supermetric configure -h supermetric-name --policies all policy- name[,{,policy-name}] --check (true false) --ht (true false) --htcriticality level-name --dtabove (true false) --dtbelow (true false) --thresholds threshold- def[,{,threshold-def}]</pre>
supermetric delete	Permanently delete a super metric.	<pre>supermetric delete -h supermetric-name</pre>

attribute Command Operations

You use the `attribute` command to configure properties of a specific metric in one or more packages. The metric is the object attribute.

The `attribute` command uses the following syntax.

```
attribute configure -h | adapterkind-key:resourcekind-key attribute-key
--packages all|package-name[,{,package-name}] --check (true|false)
--ht (true|false) --htcriticality level-name
--dtabove (true|false) --dtbelow (true|false)
--thresholds threshold-def[,{,threshold-def}]
```

reskind Command Operations for Object Types

You use the `reskind` command to configure the default settings in your object type as defined by the `ResourceKind` model element. The command sets the default attribute or supermetric package, enables or disables dynamic thresholds, and enables or disables early warning smart alerts.

The `reskind` command uses the following syntax.

```
reskind configure -h | adapterkind-key:resourcekind-key
--package package-name --smpackage smpackagename
--dt (true|false) --smartalert (true|false)
```

report Command Operations

You use the `report` command to import, export, configure, and delete super metrics.

The `report` command uses the following syntax.

```
report -h | import|export|delete [parameters]
```


Table 10-4. report Command Options

Command Name	Description	Syntax
report import	Import a report definition from a file.	<code>report import -h input-file [--force]</code>
report export	Export one or more report definitions to a file.	<code>report export -h all report-name[{{,report-name}}] [output-dir]</code>
report delete	Permanently delete one or more report definitions.	<code>report delete -h all report-name[{{,report-name}}]</code>

view Command Operations

You use the view command to import, export, or delete view definitions.

The view command uses the following syntax.

```
view -h | import|export|delete [parameters]
```

Table 10-5. view Command Operations

Command Name	Description	Syntax
view import	Import a view definition from a file.	<code>view import -h input-file [--force]</code>
view export	Export one or more view definitions to a file.	<code>view export -h all view-name[{{,view-name}}] [output-dir]</code>
view delete	Permanently delete one or more view definitions.	<code>view delete -h all view-name[{{,view-name}}]</code>

file Command Operations

You use the file command to import, export, list, or delete database files. The command operates on metric, text widget, and topology widget files.

The file command uses the following syntax.

```
file -h | import|export|delete|list [parameters]
```

Table 10-6. file Command Operations

Command Name	Description	Syntax
file import	Import a metric or widget from a file.	<pre>file import -h reskndmetric textwidget topowidget input-file [--title title] [--force]</pre>
file export	Export one or more metrics or text widgets, or export the topology widget to a file.	<pre>file export -h reskndmetric textwidget topowidget all title[{,title}] [output-dir]</pre>
file delete	Permanently delete a metric or a widget.	<pre>file delete -h reskndmetric textwidget topowidget all title[{,title}]</pre>
file list	List all metric or a widget files.	<pre>file list -h reskndmetric textwidget topowidget</pre>