vRealize Operations Manager Configuration Guide

vmware®

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

https://docs.vmware.com/

Contents

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 configure, monitor, manage, and maintain the objects in your environment.

For users who want to configure vRealize Operations Manager programmatically, the 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

<div style="text-align:right">1</div>

You can extend the monitoring capabilities of vRealize Operations Manager by installing and configuring management packs 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 management pack might be only a connection to a data source, or it might include predefined dashboards, widgets, alerts, and views.

Solutions can include cloud accounts, other accounts, dashboards, reports, alerts, and other content. The cloud accounts and other accounts comprise of adapters and using which vRealize Operations Manager manages communication and integration with other products, applications, and functions. When a management pack is installed and the adapters are configured, you can use the vRealize Operations Manager analytics and alerting tools to manage the objects in your environment.

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.

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

vRealize Operations Manager includes management packs that are pre-installed. These solutions are installed when you install vRealize Operations Manager and cannot be deactivated. The management packs are as follows:

- VMware vSphere

- VMware vRealize Log Insight

- VMware vRealize Assessments

- VMware vSAN

- vRealize Operations Service Discovery Management Pack

- VMware vRealize Automation 8.x

- VMware Management Pack for AWS

- VMware Management Pack for Microsoft Azure

- VMware vRealize Operations Management Pack for NSX-T

- VMware vRealize Operations Management Pack for VMware Cloud on AWS

- VMware vRealize Network Insight

vRealize Operations Manager also includes management packs that are bundled with vRealize Operations Manager, but not activated. You can activate these management packs from the **Repository** page. The management packs are as follows:

- Operating Systems/Remote Service Monitoring

- VMware vRealize Application Management Pack

- VMware vRealize Automation 7.x

- VMware vRealize Compliance Pack for PCI

- VMware vRealize Compliance Pack for ISO

- VMware vRealize Compliance Pack for HIPAA

- VMware vRealize Compliance Pack for FISMA

- VMware vRealize Compliance Pack for CIS

- VMware vRealize Compliance Pack for DISA

- VMware vRealize Ping

Upgrade Considerations

The management packs bundled with vRealize Operations Manager are reinstalled if vRealize Operations Manager is upgraded. If there is a fresh deployment of vRealize Operations Manager, only VMware vSphere and vRealize Optimization Assessments are installed and activated, all other management packs are pre-bundled and require activation for use.

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.

This chapter includes the following topics:

- Solutions Repository

- Managing Solutions in vRealize Operations Manager

- Managing Solution Credentials

- Managing Collector Groups

- VMware vSphere Solution in vRealize Operations Manager

- VMware Cloud on AWS

- AWS

- Microsoft Azure

- Application Monitoring

- Service Discovery

- Log Insight

- Business Management

- vRealize Automation 7.x

- vRealize Automation 8.X

- vSAN

- vRealize Network Insight

- End Point Operations Management Solution in vRealize Operations Manager

- NSX-T

Solutions Repository

You can activate or deactivate native management packs and add or upgrade other management packs from the **Repository** page.

Where You Find the Repository Page

In the menu, click **Administration**. From the left pane, select **Solutions > Repository**.

Table 1-1. Repository Page Options

Options	Descriptions
VMware Native Management Packs	
Name	Name of the solution.

Table 1-1. Repository Page Options (continued)

Options	Descriptions
Activate	Installs the native management pack. You can configure cloud management packs after activation from **Solutions > Cloud Accounts**. You can configure all other management packs after activation from **Solutions > Other Accounts**. The activation starts only if all the cluster's nodes are accessible. **Note** Pre-Installed management packs are activated by default. You can configure them from the **Cloud Accounts** or the **Other Accounts** page as applicable. Click **Add Account** configure the solutions.
Deactivate	Uninstalls the management pack. **Note** Pre-installed management packs cannot be deactivated.
Status	Indicates whether the management pack has been configured or not. A green tick symbolizes that the management pack has been successfully installed. If configured, you can view the number of accounts associated to it. To view or edit the accounts, click the account link to navigate to the accounts page associated to the management pack.
Provided By	Name of the vendor or manufacturer who created the solution.
Version	Version and build number identifiers of the solution.
View Content	Displays the list of content that has been deployed using the management pack.
Reset Default Content	This option is only available for the VMware vSphere solution. After you update your instance of vRealize Operations Manager and select the option to overwrite, alert definitions and symptom definitions, you must overwrite your existing compliance alert definitions. Reset Default Content ensures that compliance standards are current for your vSphere 6.0 and 5.5 objects. The alert definitions and symptom definitions now include the compliance standards for both vSphere 6.0 and 5.5. When you upgrade your current version of vRealize Operations Manager, you must select this option to overwrite alert definitions and symptom definitions. If you do not overwrite alert and symptom definitions, compliance rules use a mixture of new and outdated definitions.

Table 1-1. Repository Page Options (continued)

Options	Descriptions
Other Management Packs	
Add/Upgrade	You can add a management pack. For details, see Adding Solutions.

Managing Solutions in vRealize Operations Manager

You can view, activate, and configure solutions that are already installed from the Solutions page.

How Solutions Work

Solutions can include dashboards, reports, alerts and other content, cloud accounts and other accounts. The cloud accounts and other accounts contain adapters using which vRealize Operations Manager manage the communication and integration with other products, applications, and functions.

Where You Find Solutions

In the menu, click **Administration** and in the left pane under **Solutions**, click **Repository** to view and activate/deactivate cloud and other solutions. Click **Cloud Accounts** to view and configure the cloud solutions that are already installed. Click **Other Accounts** view and configure other solutions that are already installed.

Note   The VMware vSphere solution and other native management packs are pre-installed and cannot be deactivated.

Data Collection Notifications

The **Data Collection** bell icon on the menu provides quick access to status and critical notifications related to data collections. The icon indicates whether notifications exist, and whether any of them are critical.

The list displays notifications about the data collections that are in progress, and indicates whether any of them have critical issues. The list groups the data collection notifications that are in progress into a single entry at the bottom of the list. To view the details about a collection, expand the notification.

Each notification displays the status of the last or current data collection, the associated adapter instance, and the time since the collection completed or an issue was identified. You can click a notification to open the Solutions page, where you can see further details, and manage adapter instances.

If problems occur with the data collections, vRealize Operations Manager identifies those problems during each 5-minute collection cycle.

Failed Solution Installation

If a solution installation fails, plug-ins related to the solution might appear in the Plug-ins page of vRealize Operations Manager, even though the solution is not installed and does not appear on the Solutions page. When the solution installation fails, reinstall the solution.

Manage Cloud Accounts

You can view and configure cloud solutions that are already installed and configure adapter instances from the cloud accounts page.

The Cloud Accounts page includes a toolbar of options.

Click **All Filters** and select **All** to enter your criteria or filter them according to name, collector, description, solution, or adapter.

The cloud accounts page lists the solutions that were added and configured so that vRealize Operations Manager can collect data. To add another account, click Add Account and select one of the cloud solutions. For more information see, Adding Cloud Accounts.

Table 1-2. Cloud Accounts Grid Options

Option	Description
Vertical Ellipses	Change the configuration of the solution, like stop the data collection, edit or delete the cloud account, and view the object details related to the account.
Name	Name that the vendor or manufacturer gave to the solution.
Status	Indicates the status of the solution and whether the adapter is collecting any data. If the status displays a green tick with the text OK, it means that the solution is collecting data.
Description	Typically, an indication of what the solution monitors or what data source its adapter connects to.
Identifier	Version and build number identifiers of the solution.
Licensing	Indicates that the solution requires a license.
Collector	Indicates the status of the solution. Data receiving shows that the solution is collecting data.

Manage the Other Solutions

To add and configure the other solutions, see Adding Other Accounts

Adding Cloud Accounts

You can add and configure cloud accounts associated with solutions that are provided with or that you add to vRealize Operations Manager. After you have configured the account, vRealize Operations Manager can communicate with the target system. You can access the cloud accounts page at any time to modify your adapter configurations.

On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**. Click **Add Account** and select the solution you want to manage.

To manage accounts for the vSphere solution, see Cloud Account Information - VMware vSphere Account Options.

To add and configure accounts for the Management Pack for AWS, see Add a Cloud Account for Management Pack for AWS

To add and configure accounts for the Management Pack for Microsoft Azure, see Add a Cloud Account for the Management Pack for Microsoft Azure.

To add and configure accounts for VMware Cloud on AWS, see Configuring a VMware Cloud on AWS Instance in vRealize Operations Manager

Prerequisites

Note

- Activate the cloud account before adding and configuring cloud accounts.

- VMware vSphere solution is activated by default and cannot be deactivated.

Importing Cloud Accounts

You can import and synchronize existing cloud accounts from vRealize Automation 8.x to vRealize Operations Manager .The **Import Accounts** page lists all the cloud accounts associated with vCenter Server, Amazon AWS, and Microsoft Azure that are not managed by vRealize Operations Manager . You can select and import these accounts into vRealize Operations Manager directly with existing credentials as defined in vRealize Automation or add or edit the credentials before the import process. The **Import Accounts** option is hidden from the user until the integration with vRealized Automation 8.x is enabled from the integration page under **Administration** > **Management**.

Prerequisites

- Verify that vRealize Automation 8.x is enabled from **Administration** > **Management** > **Integrations** in vRealize Operations Manager .

- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data.

- Verify that the user has privileges of Organizational Owner and Cloud Assembly administrator set in vRealize Automation.

Procedure

1 In the menu, click **Administration** and then in the left pane, click **Cloud Accounts** > **Import Accounts**.

2 From the **Import Accounts** page, select the cloud account you want to import.

3   To override an existing credential from vRealize Automation, click the **Edit** icon next to **Edit
     Credential**.

   ■   Select the existing credential from the **Credential** drop-down menu and click **Save**.

   ■   To add a new credential, click the plus icon next to the **Credential** drop-down menu and
       enter the credential details and click **Save**.

4   Select the collector/group from the drop-down menu.

5   Click **Validate** to verify that the connection is successful.

6   Click **Import**.

Results

The imported cloud account is listed in the **Cloud Accounts** page. After the data collection for the
cloud account is complete the configuration status changes from **Warning** to **OK**.

Manage Other Accounts

You can view and configure native management packs and other solutions that are already
installed and configure adapter instances from the other accounts page.

Note   You must activate solutions before configuring them. For more information, see Solutions
Repository

The Other Accounts page includes a toolbar of options.

Click **All Filters** and select **All** to enter your criteria or filter them according to name, collector,
description, solution, or adapter.

The other accounts page lists the solutions that were added and configured so that vRealize
Operations Manager can collect data. To add another account, click Add Account and select one
of the solutions. For more information see, Adding Other Accounts.

Table 1-3. Cloud Accounts Grid Options

Option	Description
Vertical Ellipses	Change the configuration of the solution, like stop the data collection, edit or delete the cloud account, and view the object details related to the account.
Name	Name that the vendor or manufacturer gave to the solution.
Status	Indicates the status of the solution and whether the adapter is collecting any data. If the status displays a green tick with the text OK, it means that the solution is collecting data.
Description	Typically, an indication of what the solution monitors or what data source its adapter connects to.
Identifier	Version and build number identifiers of the solution.

Table 1-3. Cloud Accounts Grid Options (continued)

Option	Description
Licensing	Indicates that the solution requires a license.
Collector	Indicates the status of the solution. Data receiving shows that the solution is collecting data.

Manage the Cloud Solutions

To add and configure the cloud accounts, see

Adding Other Accounts

You can add and configure accounts associated with other solutions that you add to vRealize Operations Manager. After you have configured the account, vRealize Operations Manager can collect data from or send data to the target system. You can access the other accounts page at any time to modify your adapter configurations.

Note

- Activate the solutions before adding and configuring other accounts.

On the menu, click **Administration** and in the left pane, click **Solutions > Other Accounts**. Click **Add Accounts** and select the solution you want to manage.

The options available depend on the selected solution.

Configuring VMware vRealize Ping Adapter Instances

In vRealize Operations Manager, you can configure the VMware vRealize Ping functionality to verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN.

Note If you have multiple adapter instances running on different collectors and both are pinging the same address, you can still get statistics from both the adapter instances for the same IP.

Procedure

1 In the menu, click **Administration**, and then from the left pane click **Solutions** > **Other Accounts** > **Add Accounts**.

2 Click the VMware vRealize Ping adapter instance.

3 Configure the VMware vRealize Ping adapter instance.

Option	Description
Name	Enter a name for the adapter instance.
Description	Enter the description of the adapter instance.
Unique Name	Specify the name for the adapter instance. You can use the name to view the metrics published for the adapter instance.

Option	Description
Address List	Specify the IP address, IP address range, and the FQDN which must be pinged.
Configuration Filename	Specify the name of the configuration file. The configuration file contains the IP addresses, Cedar information, and FQDN details as a comma-separated file.
Collectors/Groups	Select the collector from which this adapter instance must run.
Advanced Settings	To configure the advanced settings, click the drop-down menu.
Batch Circle Interval	Specify the time interval for a new ping cycle to start. The batch circle interval values must be between 0 and 300 seconds.
Number of Pings	Specify how many times you have to ping the same IP address.
Period	Specify how long you must wait before you ping the IP address again.
DNS Name Resolve Interval	Specify the time at which you must resolve the DNS name for the next cycle. By default the value is set to 30 minutes.
Packet Size	Specify the byte size of the packet when you ping.
Don't Fragment	Select **False** to fragment the packet and **True** to not fragment the packet.

4   Click **Save**.

Results

After you configure the VMware vRealize Ping adapter instance, you can view the adapter details from **Administration** > **Solutions** > **Inventory** > **VMware vRealize Ping Adapter Instance**.

Adding Solutions

Solutions are delivered as PAK files that you upload, license, and install.

How Added Solutions Work

When you add solutions, you configure adapters that manage the communication and integration between vRealize Operations Manager and other products, applications, and functionality.

Where You Add Solutions

On the menu, select **Administration** and in the left pane select **Solutions > Repository**. Click **Add/Upgrade** to install other management packs.

Add Solutions Wizard Options

The wizard includes three pages where you locate and upload a PAK file, accept the EULA and install, and review the installation.

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.

While upgrading to the latest version, you can select the **Install the PAK file even if it is already installed** and the **Reset Default Content** options.

Table 1-4. Wizard Options

Option	Description
Page 1	
Browse a Solution	Navigate to your copy of a management pack PAK file.
Upload	To prepare for installation, copy the PAK file to vRealize Operations Manager.
Install the PAK file even if it is already installed	If the PAK file was already uploaded, reload the PAK file using the current file, but leave user customizations in place. Do not overwrite or update the solution alerts, symptoms, recommendations, and policies.
Reset Default Content	If the PAK file was already uploaded, reload the PAK file using the current file, and overwrite the solution default alerts, symptoms, recommendations, and policies with newer versions provided with the current PAK file. **Note** A reset overwrites customized content. If you are upgrading vRealize Operations Manager, the best practice is to clone your customized content before you upgrade.
The PAK file is unsigned	Warning appears if the PAK file is not signed with a digital signature that VMware provides. The digital signature indicates the original developer or publisher and provides the authenticity of the management pack. If installing a PAK file from an untrusted source is a concern, check with the management pack distributor before proceeding with the installation.
Page 2	
I accept the terms of the agreement	Read and agree to the end-user license agreement. **Note** Click **Next** to install the solution. The installation starts only if all the cluster's nodes are accessible.
Page 3	
Installation Details	Review the installation progress, including the vRealize Operations Manager nodes where the adapter was installed.

Manage Integrations

vRealize Operations Manager includes a central page where you can configure and integrate your end points to communicate with the vRealize Automation Management Pack and vRealize Log Insight Management Pack.

Where You Find Integrations

On the menu, click **Administration** and in the left pane click **Management** > **Integrations**.

Table 1-5. Integration Page Options

Property	Description
Configure	Allows you to configure and integrate your adapter instance.
Edit	Allows you to edit the integrated adapter instance.
Deactivate	Removes the adapter instance and clears the objects associated with the instance from the system, including historical data and role assignments.
Pause	Stops the data collection process.
Name	Displays the name of the Integrated adapter instance.
Version	Displays the version of the integrated adapter instance.
Status	Displays warning, OK, or Not Configured state of the integrated adapter instance.

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. To edit credential settings, click **Administration** on the menu, and in the left pane, click **Management> 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.

Credentials

The credentials are the collection configuration settings, for example, user names and passwords, that the adapters use to authenticate the connection on the external data sources. Other credentials can include values such as domain names, pass phrases, or proxy credentials. You

can configure for one or more solutions to connect to data sources as you manage your changing environment.

Where You Find Credentials

On the menu, click **Administration** and in the left pane click **Management > Credentials**.

Table 1-6. Credentials Options

Option	Description
Toolbar options	Manages the selected credential. ■ **Add**. Add new credentials for an adapter type that you can later apply when configuring an adapter. ■ Click the **Vertical Ellipses** to perform any one of the following actions: 　■ **Edit**. Modify the selected credentials, usually when the user name and password require a change. The change is applied to the current adapter credentials and the data source continues to communicate with vRealize Operations Manager. 　■ **Delete** . Remove the selected credentials from vRealize Operations Manager. If you have an adapter that uses these credentials, the communication fails and you cease monitoring the objects that the adapter was configured to manage. Commonly used to delete misconfigured credentials.
Filtering options	Limits the displayed credentials based on the adapter or credential types.
Credential name	Description of user-defined name that you provide to manage the credentials. Not the account user name.
Adapter Type	Adapter type for which the credentials are configured.
Credential Type	Type of credentials associated with the adapter. Some adapters support multiple types of credentials. For example, one type might define a user name and password, and another might define a pass code and key phrase.

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.

Note 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 1-7. 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.

Collector Group Workspace

You can add, edit, or remove collector groups in vRealize Operations Manager, and rebalance your adapter instances.

Rebalancing an Adapter Instance

Rebalancing of your adapter instances is not intended to provide equally distributed adapter instances across each collector in the collector group. The rebalancing action considers the number of resources that each adapter instance collects to determine the rebalancing placement. The rebalancing happens at the adapter instance, which can result in several small adapter instances on a single collector, and a single huge adapter instance on another collector, in your vRealize Operations Manager instance.

Rebalancing your collector groups can add a significant load on the entire cluster. Moving adapter instances from one collector to another collector requires that vRealize Operations Manager stops the adapter instance and all its resources on the source collector, then starts them on the target collector.

If a collector fails to respond or loses connectivity to the cluster, vRealize Operations Manager starts automated rebalancing in the collector group. All other user-initiated manual operations on the collector, such as to stop or restart the collector manually, do not result in automated rebalancing.

If one of the collectors fails to respond, or if it loses network connectivity, vRealize Operations Manager performs automated rebalancing. In cases of automated rebalancing, to properly rebalance the collector group, you must have spare capacity on the collectors in the collector group.

Where You Manage Collector Groups

On the menu, click **Administration** and in the left pane click **Management > Collector Groups**.

Table 1-8. Collector Group Summary Grid

Options	Description

Collector Group toolbar	To manage collector groups, use the toolbar icons.
	■ Add. Add a collector group
	■ Click the **Vertical Ellipses** to perform any one of the following actions:
	  ■ **Edit**. Modify the collector group by adding or removing remote collectors.
	  ■ **Delete**. Remove the selected collector group.
	  ■ **Rebalance collector group**. Rebalance one collector group at a time. If you have permissions to manage clusters, you can rebalance the workload across the collectors and the remote collectors in the collector group. The rebalance action moves objects from one collector group to another to rebalance the number of objects on each collector in the collector group. If a disk rebalance is already in progress, the collector rebalance does not run.

Collector Group Name	The name given to the collector group when the collector group is created.
Description	Description given to the collector group when the collector group is created.
All Filters	Displays the list of collector groups in the summary grid by collector group name, description, collector name, or IP address.
Quick Filter Name	Filters the list of collector groups according to the name of the collector group entered.

Table 1-9. Collector Group Details Grid

Detail Grid Options	Description
Members	Remote collectors that are assigned to the collector group.
Name	Name given to the remote collector when the collector was created.
IP Address	IP address of the remote collector.
Status	Status of the remote collector: online or offline

Adding a Collector Group

Create a new collector group from the available remote collectors in your environment. A collector can only be added to one group at a time.

Where You Add New Collector Groups

On the menu, click **Administration** and in the left pane click **Management > Collector Groups**. Click the **Add** icon on the Collector Groups toolbar.

Add New Collector Group Workspace

Option	Description
Name	Name of the collector group.
Description	Description of the collector group.
Members	Displays a list of the available remote collectors in your vRealize Operations Manager environment together with their IP address and status. Collectors that have already been added to a collector group are not displayed in this list.
All Filters	Enables you to search the list of collectors according to the following criteria: ■ Collector Name ■ IP address ■ Status

Editing Collector Groups

Edit a collector group by adding remote collectors to the group, or removing the collectors that you no longer require be part of the group.

Where You Edit a Collector Group

On the menu, click **Administration** and in the left pane click **Management > Collector Groups**. Click the **Edit** icon on the Collector Groups toolbar.

Edit Collector Group Options

Option	Description
Name	Name given to the collector group when the collector group is created.
Description	Description given to the collector group when the collector group is created.
Members	Displays a list of the available remote collectors in your vRealize Operations Manager environment together with their IP address and status. Collectors that have been added to another collector group are not displayed in this list. Collectors that are assigned to this collector group appear with a selected check box next to the collector name.
All Filters	Enables you to filter the list of collectors according to the following criteria: ■ Collector Name ■ IP Address ■ Status

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.

Configure the vSphere Solution to connect vRealize Operations Manager to one or more vCenter instances

To begin, access
Administration > Solutions > Configure

Configure and manage vCenter adapter instances in one central workplace

Enable/disable actions

Update the default monitoring policy

Add vCenter adapter instances

Configure user access so that users can run actions on objects in vCenter Server from vRealize Operations Manager

Create roles with permissions to determine who can access actions

Create user groups, and assign them action-specific roles and access to adapter 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.

To configure a vCenter Server cloud account, see Configure a vCenter Server Cloud Account in vRealize Operations Manager.

Configure a vCenter Server Cloud Account in vRealize Operations Manager

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

Note Any cloud account credentials you add are shared with other cloud account administrators and vRealize Operations Manager collector hosts. Other administrators might use these credentials to configure a new cloud account or to move a cloud account to a new host.

Prerequisites

- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data, see Privileges Required for Configuring a vCenter Adapter Instance. 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 > Cloud Accounts**.

2   On the Cloud Accounts page, click **Add Accounts**.

3   On the Accounts Type page, click **vCenter**.

4   Enter a display name and description for the cloud account.

   - Display name. Enter the name for the vCenter Server instance as you want it to appear in vRealize Operations Manager. A common practice is to include the IP address so that you can readily identify and differentiate between instances.

   - Description. Enter any additional information that helps you manage your instances.

5   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.

6   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.

Optionally, you can use alternate user credentials for actions. Enter an **Action User Name** and **Password**. If you do not enter an action user name and password, the default user specified is considered for actions.

Note   Credentials are stored in vRealize Operations Manager and can be used for one or more instances of the vCenter Server.

Note   To monitor application services and operating systems, it is recommended that you enter action credentials with guest operations privileges such as `guest operation alias modification`, `guest operation alias query`, `guest operation modifications`, `guest operation program execution`, `guest operation queries`.

7   Determine which vRealize Operations Manager collector or collector group is used to manage the cloud account. If you have only one cloud account, select **Default collector group**. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

8   The cloud account is configured to run actions on objects in the vCenter Server from vRealize Operations Manager. If you do not want to run actions, deselect **Enable** for Operational Actions.

9   Click **Validate Connection** to validate the connection with your vCenter Server instance.

10 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.

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

For information about these advanced settings, see Cloud Account Information - VMware vSphere Account Options.

12 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, see Cloud Account Information - VMware vSphere Account Options.

13   Click **Add** to save the configurations.

     The vCenter Server adapter instance gets saved and the vRealize Operations Manager
     Registration to the vCenter Server dialog box appears.

14   Use the vRealize Operations Manager Registration dialog box to review the registration
     information.

     ◆   If the vCenter Server already has a vRealize Operations Manager instance registered to
         it, you can override the existing registrations with your instance of vRealize Operations
         Manager. Click **Yes** to replace the existing registration with your vRealize Operations
         Manager instance.

     ◆   To proceed with the configuration without registering your vRealize Operations Manager,
         click **No**.

     You can register your vRealize Operations Manager instance after the cloud account is
     configured.

Results

The cloud account is added to the list. vRealize Operations Manager begins collecting metrics,
properties, and events 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

You can enable vSAN Configuration for your cloud account. For more information, see Configure
a vSAN Adapter Instance .

You can use the vCenter Server for service discovery, see Configure Service Discovery.

You can register your vRealize Operations Manager instance to a vCenter Server instance if you
have not done it while configuring the vCenter Server cloud account.

1   Click the cloud account you just created and click **Manage Registrations**.

    The Register vCenter Server dialog box appears.

2   Click the **Use collection credentials** check box.

    ▪   Click **Unregister** to remove any existing registrations.

    ▪   Click **Register** to register your instance of vRealize Operations Manager to the vCenter
        Server. If the vCenter Server already has a vRealize Operations Manager registered to it,
        click **Unregister** to remove the existing registration and then click **Register**.

Privileges Required for Configuring a vCenter Adapter Instance

To configure your vCenter Adapter instance in vRealize Operations Manager, you need sufficient
privileges to monitor and collect data and to perform vCenter Server actions. You can configure
these permissions as a single role in vCenter Server to be used by a single service account or
configure them as two independent roles for two separate service accounts.

The vCenter Adapter instance monitors and collects data from vCenter Server and the vCenter Action adapter performs some actions in vCenter Server. So, for monitoring or collecting vCenter Server inventory and their metrics and properties, the vCenter Adapter instance needs credentials with the following privileges enabled in vCenter Server.

Table 1-10. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection

Task	Privilege
Property Collection	**System > Anonymous**
	Note When you add a custom role and do not assign any privileges to it, the role is created as a Read Only role with three system-defined privileges: **System.Anonymous**, **System.View**, and **System.Read**. See, Using Roles to Assign Privileges .
Objects Discovery Events Collection	**Profile-Driven Storage > View** **Storage views > View** **Profile-Driven Storage > Profile-Driven Storage View** **Datastore > Browse Datastore** **System > View**
	Note This permission is provided with the Read-Only role.
Performance Metrics Collection	**Performance > Modify intervals** **System > Read**
	Note This permission is provided with the Read-Only role.
Service Discovery	**Virtual Machine > Guest Operations > Guest Operation alias modification** **Virtual Machine > Guest Operations > Guest Operation alias query** **Virtual Machine > Guest Operations > Guest Operation modifications** **Virtual Machine > Guest Operations > Guest Operation program execution** **Virtual Machine > Guest Operations > Guest Operation queries**
Tag Collection	**Global > Global tag** **Global > Global health** **Global > Manage custom attributes**
	Note This privilege is required only if the tags are associated with custom attributes.
	Global > System tag **Global > Set custom attribute**
Monitor the Namespace Resource Pool or objects in the Resource Pool.	The account for the adapter instance also needs to be a member of `Administrators@vsphere.local` on the vCenter Server.

Table 1-11. Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions

Task	Privilege
Set CPU Count for VM	**Virtual Machine > Configuration > Change CPU Count**
Set CPU Resources for VM	**Virtual Machine > Configuration > Change Resource**
Set Memory for VM	**Virtual Machine > Configuration > Change Memory**
Set Memory Resources for VM	**Virtual Machine > Configuration > Change Resource**
Delete Idle VM	**Virtual machine > Edit Inventory > Remove**
Delete Powered Off VM	**Virtual machine > Edit Inventory > Remove**
Create Snapshot for VM	**Virtual Machine > Snapshot Management > Create Snapshot**
Delete Unused Snapshots for Datastore	**Virtual Machine > Snapshot Management > Remove Snapshot**
Delete Unused Snapshot for VM	**Virtual Machine > Snapshot Management > Remove Snapshot**
Power Off VM	**Virtual Machine > Interaction > Power Off**
Power On VM	**Virtual Machine > Interaction > Power On**
Shut Down Guest OS for VM	**Virtual Machine > Interaction > Power Off**
Move VM	■ **Resource > Assign Virtual Machine to Resource Pool** ■ **Resource > Migrate Powered Off Virtual Machine** ■ **Resource > Migrate Powered On Virtual Machine** ■ **Datastore > Allocate Space** **Note** Combining these four permissions allows the service account to perform Storage vMotion and regular vMotion of an object therefore allowing vRealize Operations Manager to perform the given operations.
Optimize Container	■ **Resource > Assign Virtual Machine to Resource Pool** ■ **Resource > Migrate Powered Off Virtual Machine** ■ **Resource > Migrate Powered On Virtual Machine** ■ **Datastore > Allocate Space**
Schedule Optimize Container	■ **Resource > Assign Virtual Machine to Resource Pool** ■ **Resource > Migrate Powered Off Virtual Machine** ■ **Resource > Migrate Powered On Virtual Machine** ■ **Datastore > Allocate Space**
Set DRS Automation	**Host > Inventory > Modify Cluster**
Provide data to vSphere Predictive DRS	**External stats provider > Update** **External stats provider > Register** **External stats provider > Unregister**

For more information about tasks and privileges, see Required Privileges for Common Tasks in the *vSphere Virtual Machine Administration Guide* and Defined Privileges in the *vSphere Security Guide*.

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.

Cloud Account Information - VMware vSphere Account Options

To begin monitoring your environment with vRealize Operations Manager, you configure the VMware vSphere solution. The solution includes the vCenter Server cloud account that collects data from the target vCenter Server instances.

Where You Find the Solution - VMware vSphere

On the menu, click **Administration** and in the left pane click **Solutions > Cloud Accounts**. On the **Cloud Accounts** page, click **Add Account**, and then select the **vCenter** card.

Account Information - VMware vSphere Account Options

Configure and modify cloud accounts, and define monitoring goals on the Account Information page.

Table 1-12. Advanced Settings Options

Option	Description
Advanced Settings	Provides options related to designating specific collectors to manage this cloud account, managing object discovery and change events.
Auto Discovery	Determines whether new objects added to the monitored system are discovered and added to vRealize Operations Manager after the initial configuration of the cloud account. ■ If the value is true, vRealize Operations Manager collects the information about any new objects that are added to the monitored system after the initial configuration. For example, if you add more hosts and virtual machines, these objects are added during the next collections cycle. This is the default value. ■ If the value is false, vRealize Operations Manager monitors only those objects that are present on the target system when you configure the cloud account.
Process Change Events	Determines whether the cloud account uses an event collector to collect and process the events generated in the vCenter Server instance. ■ If the value is true, the event collector collects and publishes events from vCenter Server. This is the default value. ■ If the value is false, the event collector does not collect and publish events.
Enable Collecting vSphere Distributed Switch Enable Collecting Virtual Machine Folder Enable Collecting vSphere Distributed Port Group	When set to false, reduces the collected data set by omitting collection of the associated category.
Exclude Virtual Machines from Capacity Calculations	When set to true, reduces the collected data set by omitting collection of the associated category.
Maximum Number Of Virtual Machines Collected	Reduces the collected data set by limiting the number of virtual machine collections. To omit data on virtual machines and have vRealize Operations Manager collect only host data, set the value to zero.

Table 1-12. Advanced Settings Options (continued)

Option	Description
Provide data to vSphere Predictive DRS	vSphere Predictive DRS proactively load balances a vCenter Server cluster to accommodate predictable patterns in the cluster workload.
	vRealize Operations Manager monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.
	Predictive DRS must also be enabled for the Compute Clusters managed by the vCenter Server instances monitored by vRealize Operations Manager. Refer to the *vSphere Resource Management Guide* for details on enabling Predictive DRS on a per Compute Cluster basis.
	When set to true, designates vRealize Operations Manager as a predictive data provider, and sends predicative data to the vCenter Server. You can only register a single active Predictive DRS data provider with a vCenter Server at a time.
Enable Actions	Enabling this option helps in triggering the actions that are related to vCenter.
Cloud Type	Provides an ability to identify the type of vCenter is used in vRealize Operations Manager. By default, the cloud type is set to Private Cloud.
vCenter ID	A globally unique identifier associated with the vCenter Server instance.
Disable collecting Guest File Systems with names containing	Provide comma separated list of strings. If these strings are found in any guest files system mount point name, that guest file system will not be collected.
Collection Interval (Minutes)	The interval between collection of data from the vCenter Server.
Dynamic Thresholding	This setting is enabled by default.

The Define Monitoring Goals page provides you with default policy options which determine how vRealize Operations Manager collects and analyzes data in your monitored environment. You can change the options on this page to create a default policy.

Table 1-13. Define Monitoring Goals Page Options

Option	Description
Which objects do you want to be alerted on in your environment?	Specify the type of objects that receive alerts. vRealize Operations Manager can alert on all infrastructure objects excluding virtual machines, only virtual machines, or all.
Which types of alerts do you want to enable?	You can enable vRealize Operations Manager to trigger Health, Risk, and Efficiency alerts on your objects.
Enable vSphere Security Configuration Guide Alerts	Security Configuration Guides provide a prescriptive guidance for customers on how to operate VMware vSphere in a secure manner. Enabling this option automatically assesses your environment against the vSphere Security Configuration Guide.

You can find the vSphere Hardening Guides at http://www.vmware.com/security/hardening-guides.html.

Click **Save Settings** to finish configuration of the solution.

VMware Cloud on AWS

VMware Cloud on AWS provides the infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing private cloud like operating environment.

Configuring a VMware Cloud on AWS Instance in vRealize Operations Manager

To manage your VMware Cloud on AWS instances in vRealize Operations Manager, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS.

Prerequisites

- To configure the VMware Cloud on AWS Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.

- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.

- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

Procedure

1 On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**.

2 On the Cloud Accounts page, click **Add Accounts**.

3 On the Accounts Type page, click **VMware Cloud on AWS**.

4 Enter a display name and description for the cloud account.

- Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in vRealize Operations Manager.

- Description. Enter any additional information that helps you manage your instances.

5 To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.

- CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating CSP API Token.

- Proxy Host

- Proxy Port

- Proxy username

- Proxy Password

- Proxy Domain

6 Determine which vRealize Operations Manager collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

Note Ensure that you have Internet connectivity for the collectors to work.

7 Organization ID. Click **Get Organization** to auto fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings > View Organization**.

8 Click **Validate Connection** to validate the connection.

9 You can monitor the costs of running your VMware Cloud on AWS infrastructure by bringing in the billing from VMware Cloud on AWS to vRealize Operations Manager. To do so, enable the costing option in **Advanced Settings**.

10 Click **Save**.

The page to configure the SDDC in VMware Cloud on AWS appears.

11 Click **Configure**.

12 Configure the vCenter adapter:

a Click the **Add** icon, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.

- User Name. The vCenter user name.

- Password. The vCenter password configured for that vCenter user name.

b Select the required collector group.

c Click **Next**.

13 Configure the vSAN adapter.

a Enable the **vSAN configuration** option.

b Click **Validate Connection** to validate the connection.

c Click **Next**.

14 Configure the NSX-T adapter.

a Enable the **NSX-T configuration** option.

b Click **Validate Connection** to validate the connection.

c Click **Next**.

15 Click **Save This SDDC**.

> **Note** The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information about configuring the vCenter Service Discovery. see *Configure Service Discovery*.

The VMware Cloud on AWS account, with the configured SDDC, is added to the list.

Known Limitations

Review the following list of feature limitations of VMware Cloud on AWS integration.

- Setting the Costing Enabled option to true, collects VMware Cloud on AWS bills. The cost collection depends on the invoice generation. For instances, where the invoices are not generated, contact the VMware on Cloud for AWS support.

- vRealize Operations Manager does not calculate the datacenter, cluster, or VM costs for VMware on AWS SDDCs even if the bills are being collected. Cost drivers also cannot be configured for VMware Cloud on AWS.

- The `cloudadmin@vmc.local` user in VMware Cloud has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines on VMware Cloud. Active and consumed memory utilizations continue to work in this case.

- The compliance workflows in vRealize Operations work for virtual machines running on a vCenter Server in VMware Cloud on AWS. The compliance checks for VMware management objects such Hosts, vCenter, and so on, are not available.

- Workload optimization including pDRS and host-based business intent does not work because VMware manages cluster configurations.

- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with vRealize Operations Manager 8.0 onwards. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.

- VMware Cloud does not support vRealize Operations Manager plugin.

- You cannot log in to vRealize Operations Manager using your VMware Cloud vCenter Server credentials.

Generating CSP API Token

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of Vmware Cloud on AWS.

Prerequisites

- To configure the VMware Cloud on AWS Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.

- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.

- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

Procedure

1 Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.

2 In the **My Account** page, click **API Tokens**, and then click **Generate Token**.

3 Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.

4 Click **Generate**.

5 Copy or save the generated token.

Verify that the NSX-T Adapter Instance is Connected and Collecting Data

You configured an adapter instance of NSX-T with the VMware on AWS credentials. Now you want to verify that your adapter instance can retrieve information from the NSX-T objects in your inventory.

To view the object types, in the menu, click **Administration > Inventory > Adapter Instances > NSX-T Adapter Instance > <User_Created_Instance>**.

Table 1-14. Object Types that NSX-T Discovers

Object Type	Description
NSX-T Adapter Instance	The vRealize Operations management pack for the NSX-T instance.
Logical Switch	Logical segments in the NSX-T environment.
Logical Switches	Group of the logical segments.
Firewall Section	Firewall sections in the NSX-T environment.
Firewall Sections	Group of firewall sections.
Logical Router	Logical routers in the NSX-T environment.
Logical Routers	Group of tier-0 and tier-1 logical routers.
Tier-0 Routers	Group of tier-0 logical routers.
Tier-1 Routers	Group of tier-1 logical routers.
Group	Groups in the NSX-T environment.
Management Groups	Group of management groups in the NSX-T environment.

Table 1-14. Object Types that NSX-T Discovers (continued)

Object Type	Description
Compute Groups	Group of compute groups in the NSX-T environment.
Groups	Group of both management and compute groups.

Procedure

1 In the menu, click **Administration** and then in the left pane, click **Inventory**.

2 In the list of tags, expand **Adapter Instances** and expand **NSX-T Adapter Instance**.

3 Select the adapter instance name to display the list of objects discovered by your adapter instance.

4 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.

5 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.

AWS

Install and configure the Management Pack for AWS for vRealize Operations Manager. The Management Pack for AWS is an embedded adapter with diagnostic dashboards for vRealize Operations Manager. The adapter collects metrics from Amazon Web Services (AWS).

Introduction to the Management Pack for AWS

The Management Pack for AWS is a native management pack with diagnostic dashboards for vRealize Operations Manager. The AWS adapter collects metrics from Amazon Web Services.

Supported AWS Services

The Management Pack for AWS supports the following services in vRealize Operations Manager.

Service	Object	Description
Elastic MapReduce	EMR Job Flow	Enables developers, researchers, analysts, and data scientists to easily process vast amounts of data.
Elastic Load Balancing	Classic Load Balancer	Provides basic load balancing across multiple Amazon EC2 instances and operates at both the request level and connection level. Classic load balancer is intended for applications that are built within the EC2-Classic network.

Service	Object	Description
	Application Load Balancer	Best suited for load balancing of HTTP and HTTPS traffic, this balancer provides advanced request routing targeted at the delivery of modern application architectures, including microservices and containers.
	Network Load Balancer	Best suited for load balancing of TCP traffic where extreme performance is required.
Amazon EC2	Elastic Compute Cloud	Provides resizable computing capacity in the Amazon Web Services cloud.
	Elastic IP	Elastic IP address is a static IPv4 address designed for dynamic cloud computing, which is reachable from the Internet.
	Elastic Network Interface	Provides a logical networking component in a VPC that represents a virtual network card.
	Placement group	When you run a new EC2 instance, the EC2 service attempts to place the instance in such a way that all your instances are spread out across underlying hardware to minimize correlated failures. You can use placement groups to influence the placement of a group of interdependent instances to meet the needs of your workload.
Amazon EC2 Auto Scaling Group		Web service designed to start or stop Elastic Compute Cloud instances, based on user-defined policies, schedules, and health checks.
Amazon Elastic Block Store	EBS volume	Provides block-level storage volumes for use with Amazon Elastic Compute Cloud instances.
Amazon Relational Database Service	RDS DB Instance	Provides familiar SQL databases while automatically managing administrative tasks.
Amazon ElastiCache	ElastiCache Cluster	Amazon ElastiCache allows you to seamlessly set up, run, and scale popular open-Source compatible in-memory datastores in the cloud. Build data-intensive applications or boost the performance of your existing databases by retrieving data from high throughput and low latency in-memory data stores. Amazon ElastiCache is a popular choice for real-time use cases like Caching, Session Stores, Gaming, Geospatial Services, Real-Time Analytics, and Queuing.
	ElastiCache Node	A node is the smallest building block of an Amazon ElastiCache deployment. It is a fixed-size chunk of secure, network-attached RAM. Each node runs the engine that was selected when the cluster or replication group was created or last modified. Each node has its own Domain Name Service (DNS) name and port. Multiple types of ElastiCache nodes are supported, each with varying amounts of associated memory and computational power.
Amazon Simple Queue	SQS Queue	Provides a reliable, highly scalable, hosted queue for storing messages.

Service	Object	Description
Amazon Elastic Container Registry	ECR Container Repository	Fully managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images.
Amazon Elastic Container Service	ECS Cluster	Highly scalable, high-performance container orchestration service that supports Docker containers and allows you to easily run and scale containerized applications on AWS.
Amazon Elastic Kubernetes Service	EKS Cluster	Allows you to use Kubernetes on AWS without needing to install and operate your own Kubernetes control plane.
AWS Lambda	Lambda Function	AWS Lambda lets you run code without provisioning or managing servers.
Amazon DynamoDB	DynamoDB	Fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale.
Amazon DynamoDB Accelerator (DAX)	DynamoDB Accelerator Cluster	Fully managed, highly available, in-memory cache for DynamoDB.
Amazon Redshift	Redshift Cluster	A fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools.
Amazon Virtual Private Cloud	VPC	Lets you provision a logically isolated section of the AWS Cloud where you can run AWS resources in a virtual network that you define.
	Subnet	Provides a range of IP addresses in your VPC. Use it to run the AWS resources into a specified subnet., for example, use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that will not be connected to the Internet.
	Transit Gateway	
	Security Group	A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. When you run an instance in a VPC, you can assign up to five security groups to the instance. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC can be assigned to a different set of security groups.
	NAT Gateway	Use a network address translation (NAT) gateway to enable instances in a private subnet to connect to the Internet are other AWS services, but prevent the Internet from initiating a connection with those instances.
	VPC VPN Connection	Connect your Amazon VPC to remote networks by using a VPN connection.
Amazon CloudFront	CloudFront Distribution	AmazonCloudFront is a global content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to your viewers with low latency and high transfer speeds.

Service	Object	Description
AWS Cloudformation	Cloudformation Stack	AWS CloudFormation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment.
Amazon S3	S3 Bucket	Object storage built to store and retrieve any amount of data from anywhere.
Amazon WorkSpaces	WorkSpaces	Amazon WorkSpaces is a fully managed, secure Desktop-as-a-Service (Daas) solution that runs on AWS.
Amazon Route 53	Route53 Hosted Zone	A hosted zone is a collection of records for a specified domain.
	Route53 Health Checks	To discover the availability of your EC2 instances, a load balancer periodically sends pings, attempts connections, or sends requests to test the EC2 instances.
AWS Elastic Beanstalk	Elastic Beanstalk	Provides the fastest and simplest way to get web applications up and running on AWS. Simply upload your application code and the service automatically handles all the details such as resource provisioning, load balancing, auto-scaling, and monitoring. Elastic Beanstalk is ideal if you have a PHP, Java, Python, Ruby, Node.js, .NET, Go, or Docker web application.
Amazon Elastic File System	EFS	Provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources.

Note All services are created with the following Service Descriptors:

- Account ID
- Region
- Service Type

For more information about Amazon Web Services, go to the Amazon Web Services site at http://aws.amazon.com/.

Charges for AWS Metrics

Amazon charges you for the metrics you collect. You can reduce costs by selecting only the metrics that are most helpful and filtering out those that are of less interest.

By default, the Management Pack for AWS requests data every 5 minutes. Every collection cycle makes one Cloud Watch call per metric, per object. Currently, there are 10 basic metrics for EC2 instances and 10 basic metrics for EBS volumes. Given these figures, you can estimate the costs over time.

For information about metric costs, see http://aws.amazon.com/cloudwatch/pricing/.

Based on the costs associated with running the adapter, you can take advantage of some of the features that limit the amount of data you collect from AWS.

- Turn off auto discovery and use manual discovery. Select only those objects that are critical to your system.

- Subscribe only to specific critical regions or services.

- Use allowlist and denylist filtering to select object import by name.

- Go to the default attribute package for each object. Turn off collection of metrics that are not critical for your system.

View Management Pack for AWS Objects

You can use the inventory tree to browse and select objects. The inventory tree shows a hierarchical arrangement of the Management Pack for AWS objects by region.

Procedure

1 In the left pane of vRealize Operations Manager, click the **Environment** icon.

2 In the Environment Overview, under the Inventory Trees, click **AWS Resources by Regions**.

3 To view the child objects, expand the regions and then expand the regions per account.

   Note All the account-specific objects related to a region are grouped under the region per account section.

4 To display information about the object, select an object in the inventory tree.

Configuring the Management Pack for AWS

Configure the Management Pack for AWS in vRealize Operations Manager and optionally change its properties to customize the management pack's operation.

An Amazon Web Services account has multiple types of credentials associated with the account. Sign-in credentials are used to access the Amazon Web Services Web-based console, key pairs are used to access EC2 instances, and access keys are used in the REST API that Amazon Web Services exposes.

Because the AWS adapter is based on the REST API, you must use access keys when you set up the adapter. You generate access keys from the Amazon Web Services console. You can create credentials on a per user basis. Access keys are not a username-password pair, but a generated sequence of characters.

Note While it is not required, it is recommended that you create a guest type account, which has a read-only access to Amazon Web Services, and use the access keys associated with this account. When you create a guest group with default permissions, they do not include read access to the Elastic Map Reduce (EMR) service. You must use the IAM console to add the following permission:

```
elasticmapreduce:DescribeJobFlows
```

Generate Required Access Keys

To configure the Management Pack for AWS, you must acquire an access key and secret key from the Amazon server. You can acquire these keys as an Amazon Web Services Admin user or as an Amazon Identity and Access Management (IAM) user. For the latest instructions,

Prerequisites

- Ensure that you are using Amazon Web Services.

- Ensure that you have the valid permissions and roles in Amazon Web Services.

Procedure

1 Log in to Amazon Web Services.

2 To generate access keys, see the online documentation on the https://docs.aws.amazon.com/ site.

Complete the following tasks:

- Generate access keys as an Amazon Web Services Administrator.

- Generate access keys as Amazon Web Services Identity and Access Management User.

Configuring IAM Permissions

When you set up IAM users and groups, you can stipulate which permissions the account has for API calls. The keys you use when you set up the adapter instance must have certain permissions enabled.

For each supported AWS Service, the `ReadOnlyAccess` permission is enough to collect metrics. Use the permission to create a IAM Policy for all supported services and their related services.

Log in to the AWS console and create a json similar to the following to get the list of privileges for the service:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
             "autoscaling:Describe*",
             "cloudwatch:Describe*",
             "cloudwatch:Get*",
             "cloudwatch:List*",
             "logs:Get*",
             "logs:List*",
             "logs:Describe*",
             "logs:TestMetricFilter",
             "logs:FilterLogEvents",
             "sns:Get*",
             "sns:List*"
            ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

Table 1-15. IAM Permissions

Service	Required	Permissions
Cloudwatch	Yes.	For the list of permissions, see Cloud Watch Read Only Access json.
EC2	describeRegions is required. describeInstances and describeVolumes are only required if you subscribe to the EC2 service.	For more information, see EC2 Read Only Access json.
ELB (Elastic Load Balancing)	Required if subscribing to the ELB service.	For the list of permissions, see Elastic Load Balancing Read Only Access json.

Table 1-15. IAM Permissions (continued)

Service	Required	Permissions

EMR	Required if subscribing to the EMR service.	describe*
		
		``
		{
		    "Effect": "Allow",
		    "Action": [
		
		"elasticmapreduce:Describe*
		",
		
		"elasticmapreduce:List*",
		
		"elasticmapreduce:ViewEvent
		sFromAllClustersInConsole"
		    "s3:GetObject",
		    "s3:ListAllMyBuckets",
		    "s3:ListBucket",
		    "sdb:Select",
		
		"cloudwatch:GetMetricStatis
		tics"
		  ],
		    "Resource": "*"
		}
		``

RDS	Required if subscribing to RDS service.	For the list of permissions, see RDS Read Only Access json.
ElasticCache	Required if subscribing to ElasticCache service.	For the list of permissions, see Elastic Cache Read Only Access json.
SQS	Required if subscribing to SQS service.	For the list of permissions, see SQS Read Only Access json.
Elastic Container Registry		For the list of permissions, see Elastic Container Read Only Access json.
Elastic Container Service		list*
Lambda		For the list of permissions, see Lambda Read Only Access json and refer to the AWS Lambda policy.
DynamoDB		For the list of permissions, see Dynamo DB Read Only Access json.
DAX		describe* list*
Redshift		For the list of permissions, see Redshift Read Only Access json.
Virtual Private Cloud		For the list of permissions, see VPC Read Only Access json.
Cloud Front Distribution		For the list of permissions, see Cloud Front Distribution Read Only Access json.

Table 1-15. IAM Permissions (continued)

Service	Required	Permissions
Direct Connect		For the list of permissions, see Direct Connect Read Only Access json.
VPN Connection		describe*
VPC NAT Gateway		describe*
Elastic IP		describe*
CloudformationStack		For the list of permissions, see Cloud Formation Read Only Access json.
S3		For the list of permissions, see S3 Read Only Access json.
Workspaces		describe*
Hosted Zone		list*
Health Checks		list*

Update Configuration Settings in the Properties File

The `amazonaws.properties` file provides configuration options.

Table 1-16. Amazon Web Services Property Settings

Property	Description
firstcollecthistoryhours	Determines how far in the past to collect data when the adapter starts. The default is 0, meaning no historical collection.
maxquerywindowminutes	The maximum query window for collections, in minutes. The default is 60. The adapter asks AWS for metrics for a maximum of this many minutes.
maxhoursback	The maximum number of hours back from the current time that the adapter attempts to collect. The default value is 336, or two weeks, because Cloudwatch keeps only two weeks worth of metrics.
includetransient	False by default. Set to true to allow the adapter to import known transient objects. Transient objects currently include any EMR job that is set to terminate on completion and all of the supporting cluster EC2 instances that belong to that job.

Table 1-16. Amazon Web Services Property Settings (continued)

Property	Description
threadcount	Default is 4. Controls how many threads are active while making calls to cloudwatch to get metrics. This threadcount is per region. The total number of threads is this value times the number of regions.
collecttimeout	Controls how long the adapter waits for all metric collection calls to return from AWS during a collection cycle. The value is measured in seconds. The default value is 240 seconds, which is in line with the default 5 minute collection cycle.

Tagging Groups

The Management Pack for AWS uses tagging groups. The tagging groups appear under the AWS Entity Status in the Inventory page.

Table 1-17. Tagging Groups

Group Name	Description
PoweredOn	Objects with this tag are in the running state.
PoweredOff	Objects with this tag are in the stopped state.
Transient	Objects with this tag are not expected to persist for long periods of time.
NotExisting	Objects with this tag do not exist in the Amazon Web Services system. You can use this tag to take advantage of the periodic purge feature of vRealize Operations Manager, that the `controller.properties` file on the Analytics server controls.

Add a Cloud Account for Management Pack for AWS

You can add a Management Pack for AWS cloud account instance to your vRealize Operations Manager implementation.

Prerequisites

- Obtain the Access Key and Secret Key values. See Generate Required Access Keys. These values are not the same as your log in credentials for the Amazon Web Services site.

- Determine the services for which you collect metrics. See, Supported AWS Services

- Determine the regions to which you subscribe. Amazon Web Services is divided into nine regions. The default value * includes all regions in your subscription. If you do not want to subscribe to all regions, you can specify region identifiers in the Regions text box.

Table 1-18. Amazon Web Services Regions

Region-Friendly Name	Region Identifier
US East (N. Virginia)	us-east-1
US East (Ohio)	us-east-2
US West (N. California)	us-west-1
US West (Oregon)	us-west-2
GovCloud (US)	us-gov-west-1
Asia Pacific (Tokyo)	ap-northeast-1
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Singapore	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Osaka-Local)	ap-northeast-3
Canada (Central)	ca-central-1
China (Beijing)	cn-north-1
China (Ningxia)	cn-northwest-1
EU (Frankfurt)	eu-central-1
EU (Ireland)	eu-west-1
EU (London)	eu-west-2
EU (Paris)	eu-west-3
EU (Stockholm)	eu-north-1
South America (São Paulo)	sa-east-1
AWS GovCloud (US-East)	us-gov-east-1
AWS GovCloud (US)	us-gov-west-1
Africa (Cape Town)	af-south-1
Middle East (Bahrain)	me-south-1
Asia Pacific (Hong Kong)	ap-east-1

- Determine any blocked list or allowed list filters. These filters use regular expressions to filter in or out specific objects by name. For example, an allowed list filter of `.*indows.*` allows only objects with a name including "indows". A blocked list filter of `.*indows.*` filters out all objects with that string in their name.

Procedure

1   On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**.

2   On the Cloud Accounts page, click **Add Accounts**.

3   On the Account Types page, click **AWS**.

4   Configure the instance settings.

Option	Action
Name	Enter a name for the adapter instance.
Description	Enter a description.
Credential	Add the credentials used to access the AWS environment by clicking the plus sign. ■ Enter an instance name for the credential values you are creating. This is not the name of the adapter instance, but a friendly name for the Access Key and Secret Key credential. ■ Enter your Access Key and Secret Key values. ■ Enter any required local proxy information for your network.
Collector / Group	Select the collector upon which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default has been selected for optimal data collecting.

5   Click **Test Connection** to validate the connection.

6   Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

Option	Action
Services	Select the services from which you want to capture metrics. If you want to collect metrics for specific services, then click the drop-down icon and select one or more services. For example, `Amazon CloudFormation, Amazon EC2`. If you do not select any of the services, the metrics for all the services get collected.
Regions	Select the regions you want to subscribe to. If you want to subscribe to specific regions, then click the drop-down icon and select one or more regions. For example, `US East (N. Virginia)`,`US East (Ohio)`. If you want to subscribe to all the regions, do not select any of the regions.

Option	Action
Collect Custom Metrics	Set this option to true if you want to import all the custom metrics from you AWS account.

To publish custom metrics in vRealize Operations Manager, the metrics dimension names should match the following service mappings:

Service Name	Dimension Name
dax_cluster	ClusterId
dax_node	NodeId
dynamodb	TableName
efs	FileSystemId
eks	ClusterName
elasticbeanstalk_env	EnvironmentName
redshift_node	NodeID
redshift_cluster	ClusterIdentifier
s3_bucket	BucketName
vpc_nat_gateway	NatGatewayId
vpc_vpn	VpnId
workspace	WorkspaceId
ec2_auto_scale_group	AutoScalingGroupName
cloudfront_distribution	DistributionId
direct_connect	ConnectionId
ec2_instance	InstanceId
ec2_volume	VolumeId
transit_gateway	TransitGateway
ecs_cluster	ClusterName
ecs_service	ServiceName
elasticache_cachecluster	CacheClusterId
elasticache_cachenode	CacheNodeId
ec2_load_balancer	LoadBalancerName
application_load_balancer	LoadBalancer
network_load_balancer	LoadBalancer

Option	Action	
	Service Name	**Dimension Name**
	emr_job_flow	JobFlowId
	lambda_function	FunctionName
	rds_dbinstance	DBInstanceIdentifier
	hosted_zone	HostedZoneId
	health_check	HealthCheckId
	sqs_queue	QueueName
Support Auto Discovery	Set this option to true for automatic discovery of AWS services. If you set this value to false, when you create an adapter instance you must perform a manual discovery of services.	
Allowed List Regex	Add regular expressions to allow only objects with names that fit the criteria you specify.	
Blocked List Regex	Add regular expressions to filter out objects by name.	

7 Click **Save Settings**.

What to do next

Make sure that vRealize Operations Manager is collecting data.

Where to View the Information	Information to View
Collection Status and Collection State columns in the MP for AWS Solution Details pane on the **Cloud Accounts** page.	The collection status appears approximately 10 minutes after you have configured the adapter.
Environment Overview	The objects related to AWS are added to the inventory trees.
Dashboards	Management Pack for AWS dashboards are added to vRealize Operations Manager.

Microsoft Azure

The Management Pack for Microsoft Azure is an embedded adapter with diagnostic dashboards for vRealize Operations Manager . The adapter collects metrics from Microsoft Azure.

Supported Azure Services

The Management Pack for Microsoft Azure supports the following services.

Service	Description
Azure App Service	Allows you to build and host web applications, mobile back ends, and RESTful APIs in the programming language of your choice without managing infrastructure.
Azure Application Gateway	Allows you to build secure, scalable, and highly available web front ends in Azure. It is a web traffic load balancer that allows you to manage traffic to your web applications.
Azure Cosmos DB	A globally distributed, multi-model database service for operational and analytics workloads. It offers multiuse feature by automatically scaling throughput, compute, and storage.
Azure Kubernetes Cluster	Allows you to deploy a production ready Kubernetes cluster in Azure.
Azure Load Balancer	Allows you to evenly distribute load (incoming network traffic) across a group of backend resources or servers.
Azure MySQL Server	A fully managed database as a service offering that can handle mission-critical workloads with predictable performance and dynamic scalability.
Azure Network Interface	A network interface that allows the Azure Virtual Machine to communicate with Internet, Azure, and on-premises resources.
Azure PostgreSQL Server	A fully managed database as a service offering that can handle mission-critical workloads with predictable performance, security, high availability, and dynamic scalability. It is available in two deployment options, as a single server and as a Hyperscale (Citus) cluster.
Azure Resource Group	Allows you to use your preferred social, enterprise, or local account identities to get single sign-on access to your applications and APIs.
Azure SQL Database	A fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without any user involvement.
Azure SQL Server	Allows you to use full versions of SQL Server in the cloud without having to manage any on-premises hardware. SQL Server virtual machines (VMs) also simplify licensing costs when you pay as you go.
Azure Storage Account	Offers different access tiers, which allow you to store blob object data in the most cost-effective manner.
Azure Disk	Azure-managed disks are block-level storage volumes that are managed by Azure and used with Azure Virtual Machines. Managed disks are like a physical disk in an on-premises server but, virtualized.
Azure Virtual Machine	Provides the flexibility of virtualization without having to buy and maintain the physical hardware that runs it. However, you still must maintain the VM by performing tasks, such as configuring, patching, and installing the software that runs on it.
Azure Virtual Network	A fundamental building block for your private network in Azure. Azure Virtual Network enables many types of Azure resources, such as Azure Virtual Machines (VM), to securely communicate with each other, the Internet, and on-premises networks.
Azure Virtual Network Gateway	Virtual network gateway VMs contain routing tables and run specific gateway services. These VMs are created when you create the virtual network gateway. You cannot directly configure the VMs that are part of the virtual network gateway.

Service	Description
Azure Virtual Scale Set	Allows you to create and manage a group of load balanced VMs. The number of VM instances can automatically increase or decrease in response to demand or a defined schedule. Scale sets provide high availability to your applications, and allow you to centrally manage, configure, and update many VMs.
Azure Virtual Scale Set Instance	Allows you to create and manage a group of load balanced VMs. The number of VM instances can automatically increase or decrease in response to demand or a defined schedule. Scale sets provide high availability to your applications, and allow you to centrally manage, configure, and update many VMs.

Configuring the Management Pack for Microsoft Azure

To configure the Management Pack for Microsoft Azure, you must activate it in vRealize Operations Manager and optionally change properties to customize it.

Microsoft Azure is a native management pack. You must activate the management pack if it is deactivated. For more information, see Solutions Repository.

After activating the management pack, you must create an application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure the management pack in vRealize Operations Manager .

Note

- You can install and use the management pack only with an enterprise license of vRealize Operations Manager .

- The management pack has a default time granularity based on the services that it monitors. You cannot configure this granularity against the metrics. You can increase the collection interval but you must not decrease it. The default interval is 10 minutes.

Generate a Client Secret

Create an Active Directory application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure a cloud account for the Management Pack for Microsoft Azure.

Prerequisites

- Ensure that you are using Microsoft Azure Cloud.

- Ensure that you have a valid subscription in the Microsoft Azure portal with an Active Directory integration.

Procedure

1   Log in to the Microsoft Azure portal.

2 To create an application and generate a secret for the application, follow the instructions at https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal.

Complete the following tasks:

a Create an Azure Active Directory application.

> **Note** Ensure that the API Permission is 'Microsoft Graph User.Read'.

b Under **Access Control (IAM) > Add Role Assignment**, select the role you want to assign to the application. The minimum requirement is 'Reader' or above.

c Generate a client secret for the application.

d Copy the subscription ID, directory (tenant) ID, application (client) ID, and client secret to use in your cloud account.

Add a Cloud Account for the Management Pack for Microsoft Azure

The Management Pack for Microsoft Azure is an embedded adapter, in which each adapter instance has diagnostic dashboards, and collects metrics from Microsoft Azure. You can add a cloud account to configure an adapter instance in vRealize Operations Manager .

Prerequisites

- If the Management Pack for Microsoft Azure is deactivated, activate it in vRealize Operations Manager . For more information, see Solutions Repository.

- Generate a client secret in the Microsoft Azure portal to use in this configuration. For more information, see Generate a Client Secret.

Procedure

1 On the menu, click **Administration**.

2 In the left pane, click **Solutions > Cloud Accounts**.

3 Click **Add Account** and select **Microsoft Azure**.

4 Enter the cloud account information.

Option	Action
Name	Enter a name for the adapter instance.
Description	Enter a description for the adapter instance.

5 Configure the connection.

Option	Action
Subscription ID	Enter your subscription ID for Microsoft Azure.
Directory (Tenant) ID	Enter the directory (tenant) ID for your Azure Active Directory.

Option	Action
Credential	Add the credentials used to access Microsoft Azure by clicking the plus sign. ■ Enter an instance name for the credential values you are creating. This value is not the name of the adapter instance, but a friendly name for the secret credential. ■ Enter your application ID in your Azure Active Directory. ■ Enter the client secret that you generated for your application in the Microsoft Azure portal. ■ Enter any required local proxy information for your network.
Collector/Group	Select the collector upon which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default is selected for optimal data collecting.

6   Click **Validate Connection** to test the connection.

Note   If the test connection fails, do not add the cloud account.

If you add the cloud account with a failed test connection, vRealize Operations Manager might not collect data for the adapter instance. To resolve this issue, remove the cloud account and add it again with the correct information. If you are using a proxy, ensure that the proxy connection is efficient.

7   Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

Option	Action
Services	Select the services from which you want to collect metrics. If you want to collect metrics for specific services, then click the drop-down icon and select one or more services. For example, `Azure Disk Storage`. If you do not select any of the services, then the metrics for all the services are collected.
Regions	Select the regions you want to subscribe to. If you want to subscribe to specific regions, click the drop-down icon and select one or more regions. For example, `Central US`. If you want to subscribe to all the regions, do not select any of the regions.

8   Click **Add**.

What to do next

Ensure that the vRealize Operations Manager is collecting data.

Where to View the Information	Information to View
Environment	The objects related to the adapter instance are added to the inventory trees. For more information, see View Objects for the Management Pack for Microsoft Azure. For information about the metrics collected by the adapter, see *Metrics for the Management Pack for Microsoft Azure*.
Dashboards	The dashboards for the adapter instance are added to vRealize Operations Manager . For more information, see Microsoft Azure Dashboards.

View Objects for the Management Pack for Microsoft Azure

You can use the inventory tree in vRealize Operations Manager to browse and select objects for an adapter instance of the Management Pack for Microsoft Azure. The inventory tree shows a hierarchical arrangement of the objects by cloud account and by region.

Prerequisites

Configure an adapter instance of the Management Pack for Microsoft Azure. For more information, see Add a Cloud Account for the Management Pack for Microsoft Azure.

Procedure

1   On the menu, click **Environment**.

2   In the left pane, under **Environment Overview**, expand **VMware vRealize Operations Management Pack for Microsoft Azure**.

3   Select either of the following options:

   ■   To view the objects by region, click **Azure Resources By Region**.

   ■   To view the objects by cloud account, click **Azure Resources By Subscription**.

4   To view the object information by region, region per cloud account, subregion, cloud account, or resource group, select either of the following options:

   ■   If you are viewing objects by region, select a region. You can click the **Azure Region per Subscription** tab to view the object information for the region per cloud account. You can also expand the inventory tree for each region and select a subregion.

   ■   If you are viewing objects by cloud account, select a cloud account. You can also expand the inventory tree for each cloud account and select a resource group.

5   To view information about each object, select either of the following options:

   ■   If you are viewing objects by region, expand the inventory tree for a subregion and select an object.

   ■   If you are viewing objects by cloud account, select an object under a cloud account or expand the inventory tree for a resource group and select an object.

You can expand the inventory tree for an SQL Server object and select an SQL Database object to view information about the database object.

Application Monitoring

You can monitor application services in vRealize Operations Manager. 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 vRealize Operations Manager. You can view the data in vRealize Operations Manager and share it with the team so that they can troubleshoot the application service.

Using vRealize Operations Advanced edition, you can monitor operating systems and conduct remote checks in vRealize Operations Manager. Using vRealize Operations Enterprise edition, you can conduct remote checks, monitor operating systems and applications, and run custom scripts in vRealize Operations Manager.

vRealize Operations Manager can monitor applications using the End Point Operations Management Solution and vRealize Application Remote Collector.

Note You cannot run the vRealize Application Remote Collector agent on the same VM as the End Point Operations Management agent.

Introduction

Application monitoring enables virtual infrastructure administrators and application administrators to discover applications running in provisioned guest operating systems at a scale and to collect run-time metrics of the operating system and application for monitoring and troubleshooting respective entities. The monitoring and troubleshooting workflows are enabled from vRealize Operations Manager which include the configuration of a vRealize Operations Manager account and life-cycle management of the agents on the virtual machines.

vRealize Application Remote Collector is delivered as a standalone Photon OS OVA file. You must deploy the OVA file using a vSphere client. The OVA is available for download from vRealize Operations Manager after you log in.

The following 23 application services are supported.

Table 1-19.

Application Service	Support
Active Directory	vRealize Operations Manager
Active MQ	vRealize Operations Manager
Apache HTTPD	vRealize Operations Manager
Cassandra Database	vRealize Operations Manager
Hyper-V	vRealize Operations Manager
Java	vRealize Operations Manager
JBoss	vRealize Operations Manager
MongoDB	vRealize Operations Manager
MS Exchange	vRealize Operations Manager
MS IIS	vRealize Operations Manager
MS SQL	vRealize Operations Manager
MySQL	vRealize Operations Manager
NTPD	vRealize Operations Manager
Nginx	vRealize Operations Manager
Oracle Database	vRealize Operations Manager
Pivotal Server	vRealize Operations Manager
Postgres	vRealize Operations Manager
RabbitMQ	vRealize Operations Manager
Riak	vRealize Operations Manager
Sharepoint	vRealize Operations Manager
Tomcat	vRealize Operations Manager
Weblogic	vRealize Operations Manager
Websphere	vRealize Operations Manager

Supported Platforms

vRealize Operations Manager supports monitoring for the following platforms and app combinations with API support.

Platforms Supported by vRealize Operations Manager for Application Monitoring

Platform	Version	Architecture	Application
Red Hat Enterprise Linux	7.x 8.x	64-bit	OS Metrics and all supported applications.
CentOS	7.x	64-bit	OS Metrics and all supported applications.
Windows	Windows Server 2019 Windows Server 2016 Windows 2012 Windows Server 2012 R2	64-bit	OS Metrics and all supported applications.
SUSE Linux Enterprise Server	12.x 15.x	64-bit	OS Metrics and all supported applications.
Oracle Linux	7.x 8.x	64-bit	OS Metrics and all supported applications.
Ubuntu	18.04 LTS 16.04 LTS	64-bit	OS Metrics and all supported applications.
VMware Photon Linux	1.0 2.0 3.0	64-bit	Only OS metrics monitoring supported vRealize Application Remote Collector 8.2 runs on Photon 1.0. vRealize Application Remote Collector 8.1 runs on Photon 1.0 and vRealize Application Remote Collector 7.5 runs on Photon 1.0 Site Recovery Manager 8.2 runs on Photon 2.0 vSphere- vSphere 6.7 & 6.5 runs on Photon OS 1.0 VMware vSAN 6.7 & VMware vSAN 6.5 runs on Photon OS 1.0 Unified Access Gateway 3.7 runs on Photon 3.0 & 3.6 runs on Photon 2.0.

Sizing Reference Data

The sizing reference data helps you select a deployment configuration during the deployment of the OVA file. VMware expects vRealize Application Remote Collector sizing information to evolve, and maintains Knowledge Base articles so that sizing calculations can be adjusted to adapt to usage data and changes in versions of vRealize Operations Manager.

For more information, see the Knowledge Base article 2093783.

Supported Versions of Application Services

The application service versions which have been validated to work for application monitoring are listed here.

Application Versions Validated to Work for Application Monitoring

Application Name	Versions Validated in the Lab
Active MQ	5.15.x and 5.16.0
Apache httpd	2.4.38 2.4.39 2.4.23 2.4.6 2.2.15
Clickhouse	20.3.12.112
Java	N/A
JBoss	7.1.1 13.0 20.0.1
MongoDB	4.0.8 4.0.1 3.0.15 3.4.19
MS Exchange	MS 2016 - 15.1
MS IIS	Windows Server 2019 : 10.0.17763.1 Windows Server 2016 : 10.0.14393.0 Windows Server 2012R2 : 8.5.9600.16384 Windows Server 2012 : 8.0.9200.16384
MS SQL	Microsoft SQL Server 2014 Microsoft SQL Server 2012 Microsoft SQL Server 2017 Microsoft SQL Server 2019
My-SQL	8.0.15 5.6.35
Nginx	1.12.2
Pivotal TC server	3.2.x (3.2.8 , 3.2.14 & 3.2.13)
Postgres	11.2 10.0 9.2.23
RabbitMQ	3.6.x (3.6.15 & 3.6.10)
Redis	5:4.0.9-1ubuntu0.2

Application Name	Versions Validated in the Lab
Riak	2.1.4 2.2.3
SharePoint	2013
Apache Tomcat	9.0.17 9.0.22 8.0.33 7.0.92
Weblogic	12.2.1.3.0
Websphere	9.0 8.5.5
NTP	4.2.8p10 4.2.6p5
Active Directory	2016 2019
Hyper-V	10.0.17763.1
Cassandra Database	3.11.6 3.11.7
Oracle Database	12c 11c
Velocloud	4.0.0

Supported Versions of vCenter Server and VMware Cloud on AWS

Refer to the VMware Product Interoperability Matrix for information about the versions of vCenter Server and VMware Cloud on AWS supported for application monitoring.

Steps to Monitor Applications

You can monitor and collect metrics for your application services and operating systems.

The following flowchart describes how you can set up vRealize Application Remote Collector and vRealize Operations Manager for application monitoring.
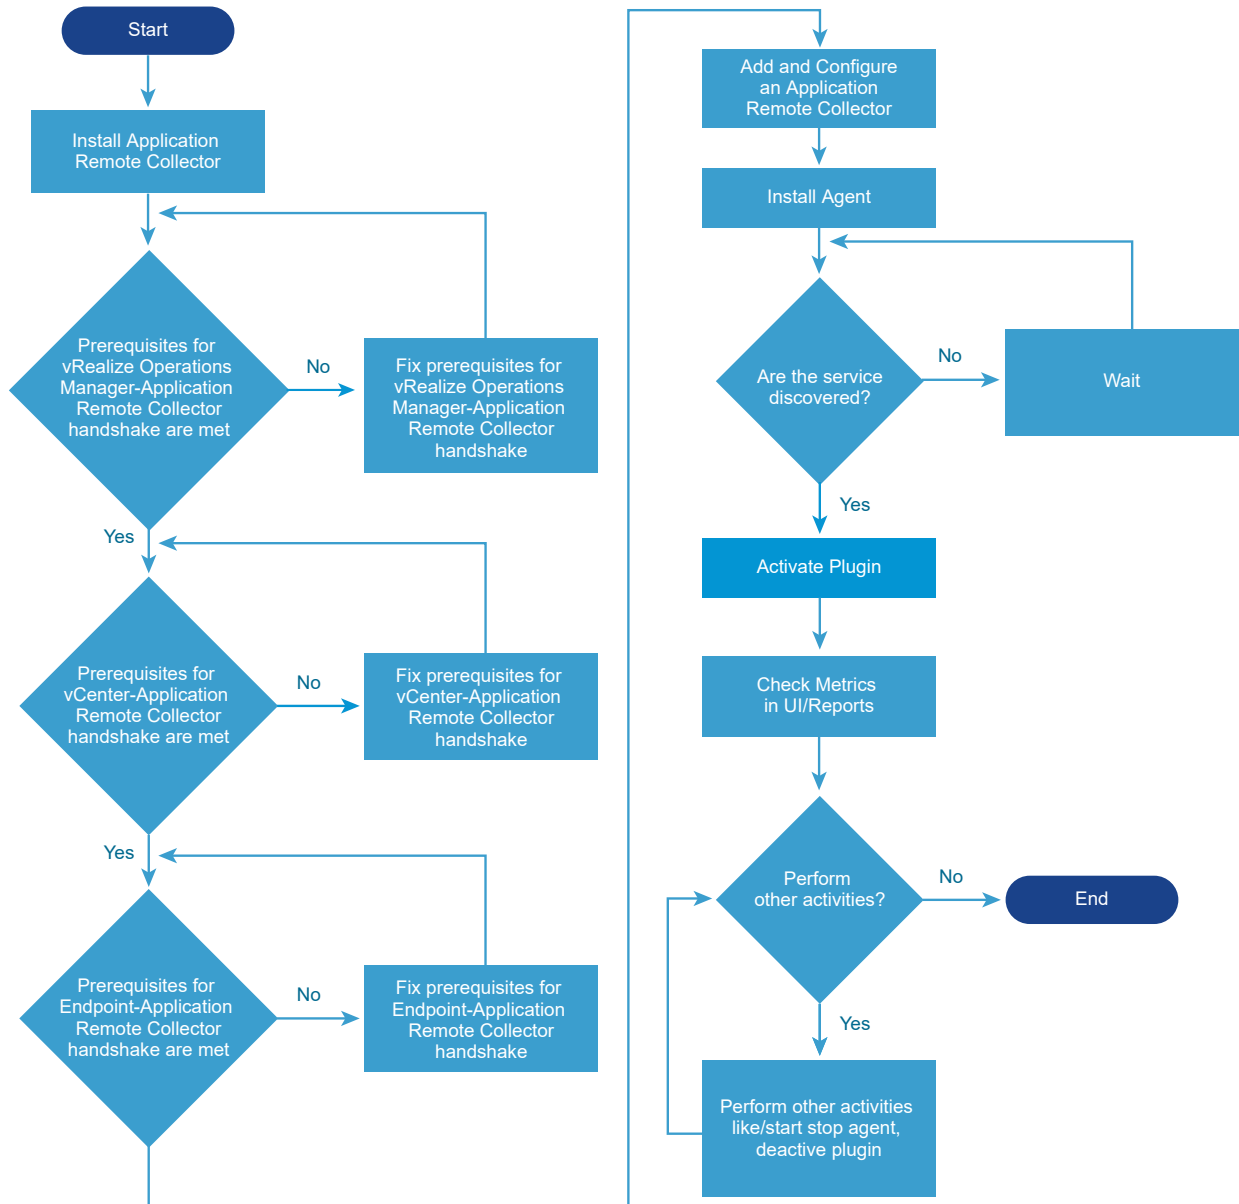
Follow these steps to monitor applications.

1   Download and deploy vRealize Application Remote Collector by clicking the **Download** icon in the **Application Remote Collector** page.

    For information about deploying vRealize Application Remote Collector, see Deploy vRealize Application Remote Collector .

2   Complete all the prerequisites.

    For more, see Prerequisites.

3   Add and configure an application remote collector.

    For information about configuring vRealize Application Remote Collector, see Application Remote Collector Page and Add and Configure an Application Remote Collector.

4   Install agents on selected VMs.

For more information, see Install an Agent from the UI.

5   Activate an application service.

For more information, see Activate an Application Service

6   View the summary of application services and operating systems discovered in vRealize Operations Manager.

For more information about monitoring your applications in vRealize Operations Manager, see Summary of Discovered and Supported Operating Systems and Application Services .

Deploy or Upgrade vRealize Application Remote Collector

Deploy vRealize Application Remote Collector

Use a vSphere client to deploy vRealize Application Remote Collector. You can deploy the vRealize Application Remote Collector OVA template from a file.

Prerequisites

You can download the vRealize Application Remote Collector OVA file after you log in to vRealize Operations Manager. Download vRealize Application Remote Collector OVA file by clicking the **Download** icon in the **Configure Application Remote Collector** page.

Note   Deployment of vRealize Application Remote Collector using vCloud Director is not supported.

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   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 vRealize Application Remote Collector OVA file, click **Local file** and browse to the location of the file and select it.

3 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.

4 On the **Select a resource** page, select a resource where to run the deployed VM template, and click **Next**.

5 On the **Review details** page, verify the OVF or OVA template details and click **Next**.

Option	Description
Product	vRealize Application Remote Collector.
Version	Version number of the vRealize Application Remote Collector.
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.

6 On the **Accept license agreements** page, click **Accept** and then **Next**.

7 In the **Select configuration** page, select the size of the deployment.

8 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.

9 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.

10 In the **Customize template** page, provide inputs to configure the vRealize Application Remote Collector deployment. It is mandatory to give these details.

Configuration	Description
API Admin User's Password	Enter a password for the vRealize Application Remote Collector API admin. The username is admin@ucp.local. This password should be used when configuring this instance of vRealize Application Remote Collector in vRealize Operations Manager. Blank spaces before or after the password are ignored and are not considered to be part of the password. Do not add the colon character : in the password.
Networking Properties	Verify the networking properties.

11 On the **Ready to complete** page, review the page and click **Finish**.

12 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

13 Change the root user password.

Note To reset the root user password, see the KB article: 2001476

14 Enable the sshd service to access the virtual machine through ssh.

What to do next

- Configure NTP Settings.

- Ensure the prerequisites for handshake with vRealize Operations Manager and vCenter Server are met.

- Log in to vRealize Operations Manager and configure application monitoring.

Configure Network Time Protocol Settings

After you install or upgrade to the latest version of vRealize Application Remote Collector, you must set up accurate timekeeping as part of the deployment. If the time settings between vRealize Application Remote Collector and vRealize Operations Manager are not synchronized, you 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 vRealize Application Remote Collector appliance 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 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
```

Upgrade vRealize Application Remote Collector

Follow the recommended upgrade flow if you have version vRealize Operations Manager prior to version 8.2, and version 8.1 of vRealize Application Remote Collector installed. Version 8.1 of vRealize Application Remote Collector is compatible with version 8.1 of vRealize Operations Manager only. Prepare for downtime during the vRealize Application Remote Collector upgrade process. There will be no flow of metrics from the VMs until the upgrade process finishes. After you upgrade vRealize Application Remote Collector, you must update the agents in the endpoints.

Recommended Upgrade Flow

- Upgrade vRealize Operations Manager from version 8.1 to version 8.2.

- Upgrade vRealize Application Remote Collector to version 8.2.

Upgrade an Existing Installation From the VAMI Portal
You must upgrade an existing installation of vRealize Application Remote Collector to ensure enhanced compatibility with vRealize Operations Manager. You must log in to your existing vRealize Application Remote Collector VAMI portal to perform the upgrade.

Prerequisites

You must have the root credentials to log in to the VAMI portal before you perform the upgrade.

Procedure

1 Log in to VAMI using root credentials. The URL to log in to VAMI is:

```
https://<IP>:5480
```

2 Click the **Update** tab.

3 Click the **Status** tab, click **Actions > Check Updates**.

4 Click **Install Updates**.

5 After the updates have installed, click **Reboot** in the **System** tab.

Results

vRealize Application Remote Collector is successfully upgraded. You can check the version number in **Update** tab under **Status** in VAMI.

What to do next

■ Update the endpoint agents to discover new services. For more information, see Additional Operations from the Manage Agents Tab.

■ To access the virtual machine appliance through ssh, start the sshd service.

■ Perform the post-installation tasks.

Upgrade an Existing Installation from an ISO File

You must upgrade an existing installation of vRealize Application Remote Collector to ensure enhanced compatibility with vRealize Operations Manager. If your deployment is behind a firewall and the VAMI portal cannot check for updates via the Internet, you can use the vRealize Application Remote Collector upgrade ISO file. You must have access to the Internet to download the vRealize Application Remote Collector upgrade ISO file.

Prerequisites

■ Download the vRealize Application Remote Collector upgrade ISO file called **VMware vRealize Application Remote Collector 8.2.0 (ISO)** from the official VMware download location.

■ You must have the root credentials to log in to the VAMI portal before you perform the upgrade.

Procedure

1 Upload the vRealize Application Remote Collector upgrade ISO file to the datastore where the vRealize Application Remote Collector appliance is deployed.

2 Power off the vRealize Application Remote Collector virtual machine.

3 Mount the vRealize Application Remote Collector upgrade ISO file to the virtual machine.

4 Power on the vRealize Application Remote Collector virtual machine.

5 Log in to VAMI using root credentials. The URL to log in to VAMI is:

```
https://<IP>:5480
```

6 Click **Install Updates** under **Status > Updates**.

7 After the updates have installed, click **Reboot** in the **System** tab.

Results

vRealize Application Remote Collector is successfully upgraded. You can check the version number in **Update** tab under **Status** in VAMI.

What to do next

- Update the endpoint agents to discover new services. For more information, see Additional Operations from the Manage Agents Tab.

- To access the virtual machine appliance through ssh, start the sshd service.

- Perform the post-installation tasks.

Prerequisites

To monitor your application services and operating systems, complete all the prerequisites so that vRealize Application Remote Collector can communicate successfully with vRealize Operations Manager, vCenter Server, and the end points.

Note For the latest port information, see https://ports.vmware.com/home

Figure 1-1. Port Information and Communication with vRealize Operations Manager, vCenter Server, and the End Points
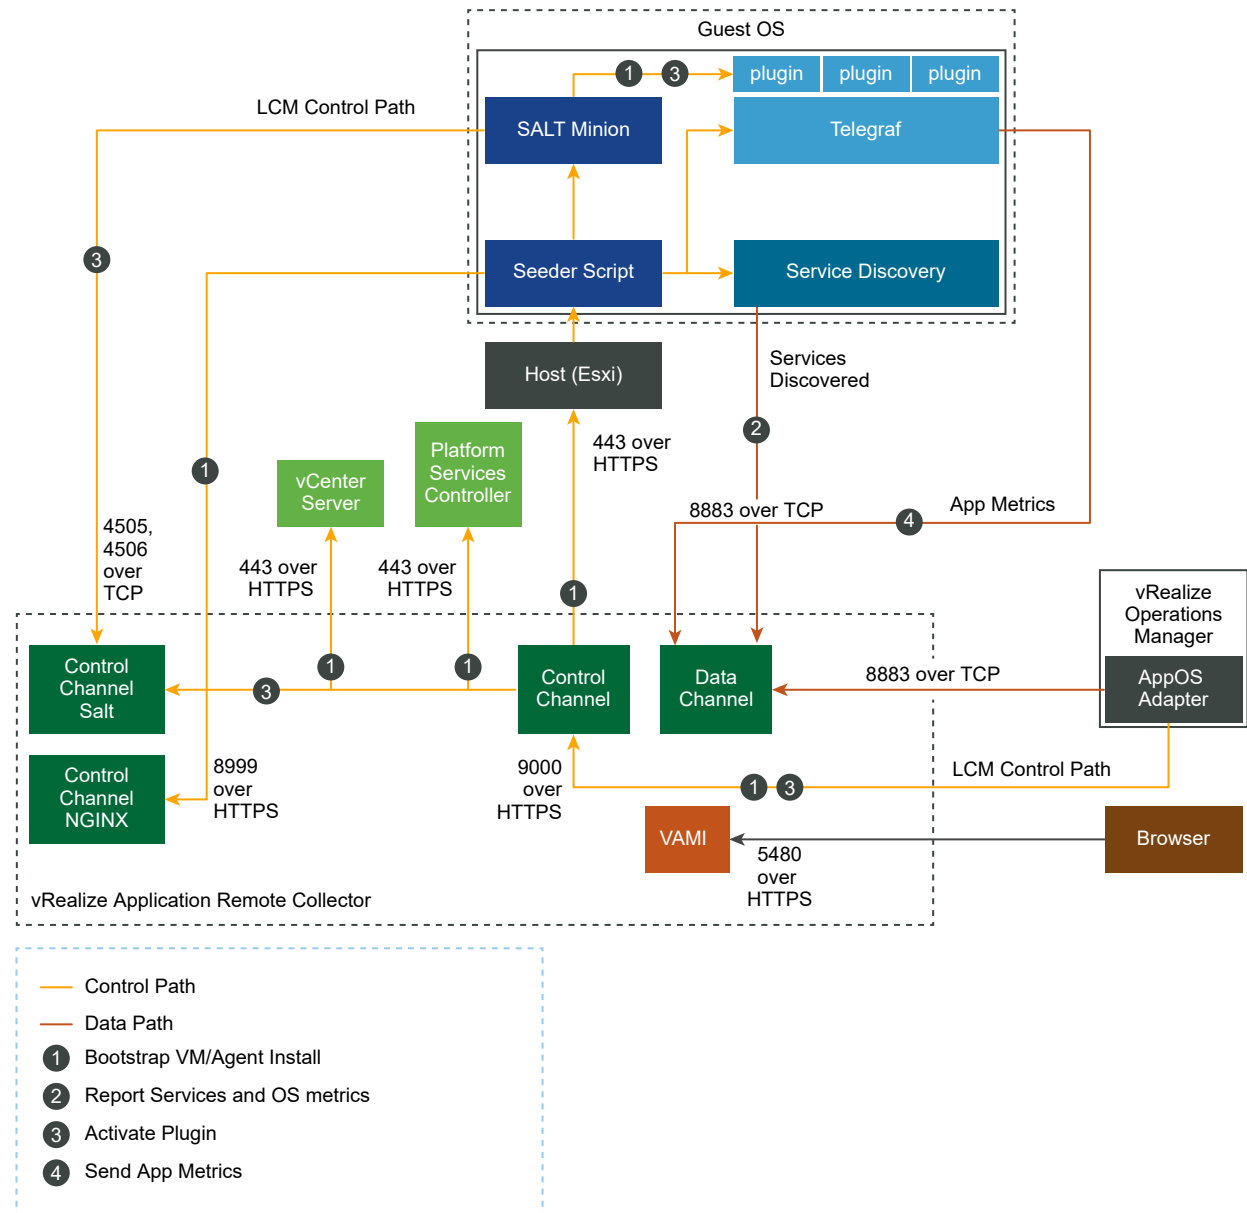
Prerequisites for Communication with vRealize Operations Manager

Ensure that you complete all the prerequisites required during the handshake of vRealize Application Remote Collector with vRealize Operations Manager.

Here are the prerequisites:

■ Verify that you have configured a vCenter adapter. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`. See Install an Agent from the UI.

- Ensure that the ports 9000 and 8883 on vRealize Application Remote Collector are reachable from vRealize Operations Manager.

- Download and deploy vRealize Application Remote Collector.

  You can download vRealize Application Remote Collector by clicking the **Download** icon in the **Configure Application Remote Collector** page.

  For information about deploying the vRealize Application Remote Collector, see Deploy vRealize Application Remote Collector .

- Ensure that the NTP settings of vRealize Operations Manager and vRealize Application Remote Collector are in sync. To configure NTP, see Configure Network Time Protocol Settings.

Prerequisites for Communication with vCenter Server

Ensure that you complete all the prerequisites required so that vRealize Application Remote Collector can communicate with vCenter Server.

- Ensure that the NTP settings of the ESXi instance that hosts the end points and vRealize Application Remote Collector are in sync.

- Port 443 in vCenter Server is accessible to vRealize Application Remote Collector.

- Port 443 in the ESXi where the workload end-points are deployed must be accessible to vRealize Application Remote Collector.

- Port 443 in Platform Services Controller is accessible to vRealize Application Remote Collector. Open this port if vCenter Server is configured with an external Platform Services Controller.

- Verify that you have configured a vCenter adapter. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have read access at the vCenter Server level and should also have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`. See Install an Agent from the UI.

Prerequisites for Communication with the End Points

Ensure that you complete the prerequisites required during the handshake of vRealize Application Remote Collector with the end points.

Here are the prerequisites:

- Ensure that the NTP settings of the ESXi instance that hosts the end points, the end points, and vRealize Application Remote Collector are in sync.

- Ensure that the end points have access to ports 8999, 4505, 4506, and 8883 on vRealize Application Remote Collector.

- Guest operation privileges are required to install agents on virtual machines. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.

- Account privilege prerequisites. See User Account Prerequisites for more details.

- End-point VM configuration requirements.

  - Linux requirements

    Commands: `/bin/bash, sudo, tar, awk, curl`

    Packages: `coreutils (chmod, chown, cat), shadow-utils (useradd, groupadd, userdel, groupdel)`

    Configure mount point on `/tmp` directory to allow script execution.

  - Windows 2012 R2 requirement

    The end point must be updated with the Universal C Runtime. Refer to the following link for more information.

  - Windows requirement

    - The Visual C++ version must be higher than 14.

    - Performance Monitors on a Windows OS VM must be enabled.

- VMware Tools must be installed and running on the VM on which you want to install the agent. For information about supported VMware Tools versions, click this link.

- If plugin activation requires the location of a file (for example, client certificates for SSL Trust) on the endpoint VM, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

  Note If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

User Account Prerequisites

There are certain user account prerequisites required for the install of agents.

Prerequisites for Windows End Points

- To install agents,

  - The user must be either an administrator, or

  - A non-administrator who belongs to the administrator group.

Prerequisites for Linux End Points

- `/tmp` mount point should be mounted with exec mount option.

- Ensure that the following lines exist in `/etc/sudoers`.

```
1.root ALL=(ALL:ALL) ALL
2.Defaults:root !requiretty
3.Defaults:arcuser !requiretty
```

(1) can be omitted if password-less sudo is already enabled for the root user. (2) and (3) can be omitted if your end point VMs are already configured to turn off `requiretty`.

For Linux end points, there are two user accounts, such as the install user and the run-time user.

Install User Prerequisites

You can use one of the following install users for Linux end points.

- root user - All privileges

- A non-root user with all privileges -

  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 *bob*, add `bob ALL=(ALL:ALL) NOPASSWD: ALL` to `/etc/sudoers`.

  To enable password-less sudo elevation access for a user group called *bobg*, add `%bobg ALL=(ALL:ALL) NOPASSWD: ALL` to `/etc/sudoers`.

- A non-root user with a specific set of privileges -

  Password-less sudo elevation access for a non-root user with access to certain commands. To enable password-less sudo elevation access for the ARC_INSTALL_USER, add the following corresponding entries to the *sudoers* file:

```
Defaults:ARC_INSTALL_USER !requiretty
Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/
cp*,/usr/bin/mkdir*,/bin/mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-
bootstrap.sh,/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh
ARC_INSTALL_USER ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS

For example,for a user bob, add the following lines to /etc/sudoers:
Defaults:bob !requiretty
Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/
cp*,/usr/bin/mkdir*,/bin/mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-
bootstrap.sh,/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh
bob ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS
```

Run-Time User Prerequisites

There are two ways in which a run-time user is created in Linux end points: automatically and manually. A run-time user has a standard name and group, which is the *arcuser* and *arcgroup* respectively. By default, the *arcuser* and *arcgroup* are created automatically. If you choose to manually create the *arcuser* and *arcgroup*, here are the prerequisites:

- Manually created *arcuser* and *arcgroup*.

Create the *arcgroup* and *arcuser* and associate the *arcgroup* as the primary group of the *arcuser*. Here are the requirements:

a   The *arcgroup* must be the primary group of the *arcuser*.

For example, the following commands can be used to create the *arcgroup* and *arcuser*:

```
groupadd arcgroup
```

```
useradd arcuser -g arcgroup -M -s /bin/false
```

b   The *arcuser* must be created with no home directory and no access to the login shell.

For example, the `etc/passwd` entry for the *arcuser* is as follows after adding *arcuser* and *arcgroup*.

```
arcuser:x:1001:1001::/home/arcuser:/bin/false
```

c   The *arcuser* must have either password-less all privileges or password-less specific set of privileges as mentioned below:

To enable password-less sudo elevation access for the run-time *arcuser*, add the following corresponding entries to the *sudoers* file.

All privileges:

```
arcuser ALL=(ALL:ALL) NOPASSWD: ALL
```

Specific set of privileges:

```
Cmnd_Alias ARC_RUN_COMMANDS=/usr/bin/systemctl * ucp-telegraf*,/bin/systemctl * ucp-
telegraf*, /usr/bin/systemctl * ucp-minion*, /bin/systemctl * ucp-minion*, /usr/bin/
systemctl * salt-minion*, /bin/sytemctl * salt-minion*, /usr/bin/netstat, /bin/
netstat, /opt/vmware/ucp/tmp/telegraf_post_install_linux.sh, /opt/vmware/ucp/bootstrap/
uaf-bootstrap.sh, /opt/vmware/ucp/uaf/runscript.sh, /opt/vmware/ucp/ucp-minion/bin/ucp-
minion.sh
arcuser ALL=(ALL) NOPASSWD: ARC_RUN_COMMANDS
```

Configure Network Time Protocol Settings

After you install or upgrade to the latest version of vRealize Application Remote Collector, you must set up accurate timekeeping as part of the deployment. If the time settings between vRealize Application Remote Collector and vRealize Operations Manager are not synchronized, you 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 vRealize Application Remote Collector appliance 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 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 Remote Collector

You can add and configure an application remote collector from the **Application Remote Collector** page to manage the life cycle of agents and application services.

To add and configure a vRealize Application Remote Collector, in the menu, click **Administration**, and then in the left pane select **Configuration > Application Remote Collector**.

Note   Time synchronization between vRealize Application Remote Collector and vRealize Operations Manager is mandatory when you add an application remote collector. If the time settings are not synchronized, you face problems such as, a failed test connection when you add an application remote collector, agent installation issues, and issues in metrics collection after the agent is installed. For more information, see Troubleshoot Agent Installation and Metric Collection Issues.

For more troubleshooting information on vRealize Application Remote Collector, see Troubleshooting the Configuration of vRealize Application Remote Collector.

Prerequisites

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

Procedure

1 To configure a vRealize Application Remote Collector, click the **Add** icon from the **Application Remote Collector** page.

2 In the **Application Remote Collector** page, enter the following details:

    a FQDN of the vRealize Application Remote Collector you have configured during the installation of vRealize Application Remote Collector.

    b You cannot modify the user name which is `admin`.

    c The API password of the vRealize Application Remote Collector you have configured during the installation of vRealize Application Remote Collector.

    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 vRealize Application Remote Collector.

    If you have mapped a vCenter Server to a vRealize Application Remote Collector, it is not displayed in the drop-down menu.

    b The vCenter Servers that are mapped to the vRealize Application Remote Collector 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 vRealize Application Remote Collector. If the mapped vCenter Server turns green, it signifies that vRealize Operations Manager can communicate with the vRealize Application Remote Collector.

    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 vRealize Application Remote Collector.

    It might take up to 5 minutes to get the status of vRealize Application Remote Collector.

    a Click **Finish**.

What to do next

Install agents on the VMs you prefer and manage the application services.

Application Remote Collector Page

The application remote collectors you add and configure are displayed in the **Application Remote Collector** page.

You can view the name of the vRealize Application Remote Collector added and the number of vCenters managed, in the **Application Remote Collector** page.

Table 1-20. Options

Options	Description
Add	You can map a vCenter Server with a vRealize Application Remote Collector as part of the configuration process. For more information, see Add and Configure an Application Remote Collector.
	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.
Edit	You can modify the vRealize Application Remote Collector configuration details or the details of the vCenter Servers that are managed.
	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.
Delete	You can delete the application remote collector. Ensure that you uninstall the agents from the VMs that are monitored before you delete the application remote collector.
Download	You can download vRealize Application Remote Collector. For information about deploying vRealize Application Remote Collector, see Deploy vRealize Application Remote Collector .

You can also view specific details from the options in the data grid.

Table 1-21. Data Grid Options

Option	Description
Name	Displays the FQDN of the vRealize Application Remote Collector.
Application Remote Collector Version	Displays the version of vRealize Application Remote Collector. A gray dot is displayed if there is a newer version of vRealize Application Remote Collector available.
vCenters Managed	Displays the number of vCenter Servers mapped to the vRealize Application Remote Collector.
Collector Server Status	Indicates the health of the vRealize Application Remote Collector.
	■ Green. Indicates that the vRealize Application Remote Collector is healthy.
	■ Red. Indicates that the vRealize Application Remote Collector is not healthy.
	Point to this cell to view a tooltip that displays the cause if the health status is red.
	The progress status is displayed when data collection has not started.

Under **Advanced Settings**, the collection interval is set to 5 minutes.

Install an Agent

You can install agents on a VM from the user interface of vRealize Operations Manager or by running a script.

Install an Agent from the UI

You must select the VMs on which you want to install the agent. If you have upgraded an existing installation of vRealize Application Remote Collector, upgrade the agents that you have previously installed.

Prerequisites

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

Procedure

1   From the **Manage Agents** tab, click the **Install** icon. 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   The **Create run time user on Linux virtual machines, with required permissions as part of agent installation** check box is selected by default. For more information, see User Account Prerequisites.

   d   Click **Next**.

4   From the **Summary** page, you can view the list of VMs on which the agent is to be deployed.

5   Click **Install Agent**. Refresh the UI to view the agents that are installed.

   On UAC disabled machines on Windows end points, the agent discovers the application services that are installed on the VMs. The application services are displayed in the **Services Discovered/Configured** column in the **Manage Agents** tab. You can view the status of agent installation from the **Agent Status** column in the **Manage Agents** tab.

   UAC Enabled Machines on Windows End Points

The bits are downloaded to the end point. You have to manually install the bits.

a   From `C:\VMware\UCP\downloads`, run a bootstrap launcher.

b   Go to `%SYSTEMDRIVE%\VMware\UCP\downloads`.

c   Open `cmd` with administrator privileges.

d   Run the `cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1` command.

e   View the results from `uaf_bootstrap.log`.

f   Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Agents** tab.

What to do next

You can manage the services on each agent.

For information about uninstalling an agent, see Uninstall an Agent.

Install/Uninstall an Agent Using a Script on a Linux Platform

You can install or uninstall an agent on a VM using a script.

Prerequisites

■   Ensure that the end point is available in vRealize Operations Manager.

■   Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

■   Ensure that the unzip package is available on the VM.

■   Ensure that the user has access permissions to the download folder.

■   Ensure that the guest IP is properly configured and is unique across vCenter Servers. If more than one VM with the same IP across vCenter Servers is monitored, the script cannot resolve and subscribe to application monitoring.

■   Ensure that the cloud account is configured for the vCenter Server to which the VM belongs. The vCenter Server must be mapped with vRealize Application Remote Collector.

■   Only IPv4 is supported at present.

Procedure

1 Log in to the VM on which you want to install/uninstall the agent and download the sample script from vRealize Application Remote Collector from the following location: `https://<ApplicationRemoteCollectorIP>:8999/downloads/salt/download.sh`.

Run one of the following commands:

```
wget --no-check-certificate https://<ApplicationRemoteCollectorIP>:8999/downloads/salt/
download.sh
curl -k "https://<ApplicationRemoteCollectorIP>:8999/downloads/salt/download.sh" --output
download.sh
```

Note Use the relevant vRealize Application Remote Collector IP address/FQDN for < ApplicationRemoteCollectorIP> in the preceding commands and location specified.

2 Make the script executable by running the following command:

```
chmod +x download.sh
```

3 To execute the script and install/uninstall the agent, run the following command:

```
./download.sh  -o <operation> -v <vrops_ip_or_fqdn> -u <vrops_user> -p <vrops_password>
[-d download_tmp_dir]

Description of arguments:
operation - Bootstrap operation. values: install, uninstall.
vrops_ip_or_fqdn - IP/FQDN of vRealize Operations Manager. This can be the address of any
vRealize Operations Manager node or VIP of vRealize Operations Manager.
vrops_user - vRealize Operations Manager user. The user should have enough permissions.
vrops_password - Password of vRealize Operations Manager.
download_tmp_dir - Temporary directory to download agent related bits. It's an optional
parameter. Default value: current directory.
```

To verify the bootstrap status, verify the `uaf-bootstrap-results file`.

If the script is successful, the agent status will be updated in the **Manage Agents** tab after one collection cycle that takes 5–10 minutes.

Note When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

Install/Uninstall an Agent Using a Script on a Windows Platform

You can install an agent on a VM using a script.

Prerequisites

▪ Ensure that the end point is available in vRealize Operations Manager.

▪ Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

- Ensure that the unzip package is available on the VM.

- Ensure that the user has access permissions to the download folder.

- Ensure that the Windows PowerShell is >= 4.0.

- Ensure that the guest IP is properly configured and is unique across vCenter Servers. If more than one VM with the same IP across vCenter Servers is monitored, the script cannot resolve and subscribe to application monitoring.

- Ensure that the cloud account is configured for the vCenter Server to which the VM belongs. The vCenter Server must be mapped with vRealize Application Remote Collector.

- Only IPv4 is supported at present.

Procedure

1   Log in to the VM on which you want to install/uninstall the agent and download the sample script from vRealize Application Remote Collector from the following location: `https://<ApplicationRemoteCollectorIP>:8999/downloads/salt/download.ps1`

Run one of the following commands:

```
Invoke-WebRequest "https://<ApplicationRemoteCollectorIP>:8999/downloads/salt/
download.ps1" -OutFile download.ps1
wget --no-check-certificate https://<ApplicationRemoteCollectorIP >:8999/downloads/salt/
download.ps1
```

Note   Use the relevant vRealize Application Remote Collector IP address/FQDN for < ApplicationRemoteCollectorIP> in the preceding commands and location specified.

2   To execute the script and install/uninstall the agent, run the following command:

```
powershell -file .\download.ps1  -o <operation> -v <vrops_ip_or_fqdn> -u <vrops_user> -p
<vrops_password> [-d download_tmp_dir]

Description of arguments:
operation - Bootstrap operation. values: install, uninstall.
vrops_ip_or_fqdn - IP/FQDN of vRealize Operations Manager. This can be the address of any
vRealize Operations Manager node or VIP of vRealize Operations Manager.
vrops_user - vRealize Operations Manager user. The user should have enough permissions.
vrops_password - Password of vRealize Operations Manager.
download_tmp_dir - Temporary directory to download agent related bits. It is an optional
parameter. Default value: current directory.
```

To verify the bootstrap status, verify the `uaf-bootstrap-results` file.

If the script is successful, the agent status will be updated in the **Manage Agents** tab after one collection cycle that takes 5–10 minutes.

Note   When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

Activate an Application Service

To monitor application services running on the target VMs, plugins must be configured in the target VMs after the agent is installed.

After you have installed the agent, you can activate plugins to monitor application services. You can also reactivate plugins that must be monitored.

Prerequisite

- If plugin activation requires the location of a file (for example, client certificates for SSL Trust) on the endpoint VM, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

  Note If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

Activate an Application Service

To monitor an application service, complete the following steps:

1   Navigate to the **Inventory > Manage Agents** tab.

2   Select the VM on which agent is already installed.

3   Select **Manage Service** icon and then from the drop-down menu select the **service name**.

4   Activate the application service from the right pane of the **Manage <service name> Agent** dialog box.

5   Click the **Add** icon in the left pane to add multiple instances of the application service.

6   Click the **Delete** icon in the left pane to delete instances of the application service.

7   Enter the details for each instance that you add and click **Save**. For configuration details of each application, see Configuring Supported Application Services.

    For more information about the status details that appear against the application services in the **Services Discovered/Configured** column, see the table called Data Grid Options in Additional Operations from the Manage Agents Tab.

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 When multiple VMs are selected, the **Manage Service** option is disabled.

For information about deactivating a service, see Deactivate an Application Service.

Configuring Supported Application Services

Twenty-three application services are supported in vRealize Operations Manager. The supported application services are listed here. Some of the application services have mandatory properties

which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

Active Directory

Active Directory is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

Active MQ

ActiveMQ is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Server URL	Yes	http://localhost:8161
User name	Yes	User name for Active MQ. Example: admin
Password	Yes	Password
Installed Path	Yes	The path on the Endpoint where Active MQ is installed. Example: For Linux VMs: `/opt/apache-activemq` For Windows VMs: `C:\apache-activemq-5.15.2`

Apache HTTPD

Apache HTTPD is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Status Page URL	Yes	http://localhost/server-status?auto
User name	No	User name for Apache HTTPD service. **Example:root**
Password	No	Password
SSL CA	No	Path to the SSL CA file on the Endpoint
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint

Name	Mandatory?	Comment
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

Cassandra Database

Cassandra database is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the application instance.
Installed Path	Yes	Valid file path.
URL	Yes	http://localhost:8778

Hyper-V

Hyper-V is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the application service.

Java

Java is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Base URL	Yes	http://localhost:8080
Installed Path	Yes	The path on the Endpoint where Java is installed. Example: For Linux VMs : /opt/vmware/ucp ; For Windows VMs : C:\VMware\UCP
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

JBoss

JBoss is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Base URL	Yes	http://localhost:8080
Installed Path	Yes	The path on the Endpoint where JBoss is installed.
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

MongoDB

MongoDB is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Port	Yes	The port where MongoDB is running. Example: 27017
Hostname	No	Optional hostname for the MongoDB Service.
Username	No	User name for MongoDB. Example: Root
Password	No	Password
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

MS Exchange

MS Exchange is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

MS IIS

MS IIS is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

MS SQL

MS SQL is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Instance	Yes	Instance name of the MS SQL server
Port	No	The port where MS SQL is running. Example: 1433
Hostname	No	Optional hostname for the MS SQL Service.
Username	Yes	User name for MS SQL. Example: Root
Password	Yes	Password

MySQL

MySQL is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Port	Yes	The port where MySQL is running. Example: 3306
User name	Yes	User name for MySQL service. Example: Root
password	Yes	Password
SSL CA	No	Path to the SSL CA file on the Endpoint
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint

Name	Mandatory?	Comment
SSL Key	No	Path to the SSL Key file on the Endpoint.
Hostname	No	Optional hostname for the MySQL Service
Databases	No	Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma separated. For example, 'database1','database2','database3'.
TLS Connection	No	Allowed values are true, false, and skip-verify.

NTPD

NTPD is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

Oracle Database

Oracle database is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the application instance.
OracleDB Username	Yes	User name for the Oracle database instance.
OracleDB Password	Yes	Password for the Oracle database instance.
OracleDB SID	Yes	SID of the Oracle database instance.

Pivotal Server

Pivotal Server is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Base URL	Yes	http://localhost:8080
Installed Path	Yes	The path on the Endpoint where Pivotal server is installed.

Name	Mandatory?	Comment
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

Postgres

Postgres is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Port	Yes	The port where PostgreSQL is running. Example: 5432
User name	Yes	User name for PostgreSQL service. Example: Root
Password	Yes	Password
SSL Connection	No	Allowed values are disable, verify-ca, verify-full.
SSL CA	No	Path to the SSL CA file on the Endpoint
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: true/false.
Hostname	No	Optional hostname for the PostgreSQL Service.
Default Database	No	The database for initiating connection with the server

Name	Mandatory?	Comment
Databases	No	Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated, for example , 'database1','database2','database3'.
Ignored Databases	No	Comma-separated list of databases that need not be monitored. Each of the database names to be excluded from monitoring must be enclosed in single quotes and the databases themselves should be comma-separated for example, 'database1','database2','database3'.

RabbitMQ

RabbitMQ is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Management Plugin URL	Yes	http://localhost:15672
User name	No	User name for RabbitMQ. Example: Guest
Password	No	Password
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.
Nodes	No	Each of the RabbitMQ data collection nodes should be in single quotes and the nodes themselves should be comma-separated. The list of nodes must be enclosed in square brackets. For example ['rabbit@node1','rabbit@node2',.....]

Riak

Riak is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Server URL	Yes	http://localhost:8098

Sharepoint

Sharepoint is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

Tomcat

Tomcat is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Base URL	Yes	http://localhost:8080
Installed Path	Yes	The path on the Endpoint where Tomcat is installed.
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

Weblogic

Weblogic is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Base URL	Yes	http://localhost:7001
Installed Path	Yes	The path on the Endpoint where WebLogic is installed.
User name	Yes	User name for WebLogic. Example: admin
Password	Yes	Password

Name	Mandatory?	Comment
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification	No	Use SSL but skip chain & host verification. Expected: True/False.

Websphere

Websphere is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
IBM Websphere Server URL	Yes	Example : http://localhost:9081
Websphere Authorization Token	Yes	To generate the token, follow the below steps: ■ Go to https://www.base64encode.org. ■ Type in the user and password created in the format: user:password ■ Click the Encode button. ■ Copy the resulting Base64 encoded string. Example: d2F2ZWZyb250OndhdmVmcm9udA==

Remote Checks

HTTP Remote Check

HTTP is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the remote check instance.
URL	Yes	http://localhost
Method	Yes	GET/POST/PUT
Proxy	No	Proxy URL: http://localhost
Response Timeout	No	Timeout for the connection in seconds. For example, 10.

Name	Mandatory?	Comment
Follow Redirects	No	True/False if redirects from the server. For example, true/false (all small values).
Body	No	HTTP request body.
Response String Match	No	Substring or regex match in the response body.
SSL CA	No	Path to the SSL CA file on the end point.
SSL Certificates	No	Path to the SSL certificate file on the end point.
SSL Key	No	Path to the SSL key file on the end point.
Skip Host & chain verification	No	Use SSL but skip chain and host verification. Expected: True/False.

ICMP Remote Check

ICMP is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the remote check instance.
FQDN/IP	Yes	Host name to send the packets. Example: *example.org*
Count	No	Number of ping packets to send per interval. For example, 1.
Ping Interval	No	Time to wait between ping packets in seconds. For example, 10.0. **Note** Follow the decimals as mentioned in the example.
Timeout	No	Timeout to wait for ping response in seconds. For example, 10.0. **Note** Follow the decimals as mentioned in the example.
Deadline	No	The total ping deadline in seconds. For example, 30.
Interface	No	Interface or source from which to send a ping.

TCP Remote Check

TCP is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the remote check instance.
Address	Yes	<hostname>:port
Send	No	The given string is sent to the TCP. It can be any string of your choice.
Expect	No	The given string is expected from the TCP. It can be any string of your choice.
Timeout	No	Timeout for the connection to the TCP server. For example, 10.
Read Timeout	No	Timeout for the response from the TCP server. For example, 10.

UDP Remote Check

UDP is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the remote check instance.
Address	Yes	<hostname>:port
Send	Yes	The given string is sent to the UDP.
Expect	Yes	The given string is expected from the UDP.
Timeout	No	Timeout for the connection to the UDP server. For example, 10.
Read Timeout	No	Timeout for the response from the UDP server. For example, 10.

Configuring Supported VeloCloud Services

Eight VeloCloud application services are supported in vRealize Operations Manager. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

VeloCloud Orchestrator

VeloCloud Orchestrator and the following services are supported in vRealize Operations Manager.

- Java Application

- VeloCloud Orchestrator

- Nginx

- Clickhouse

- Network Time Protocol

- MySQL

- Redis

In VeloCloud Orchestrator, we monitor the following services. For each of these services we display a metric which indicates the service status:

- Backend

- Portal

- Upload

VeloCloud Orchestrator details.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the VeloCloud Orchestrator instance.

Nginx

Nginx is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Status Page URL	Yes	http://127.0.0.1/nginx_status
SSL CA	No	Path to the SSL CA file on the Endpoint.
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint.
SSL Key	No	Path to the SSL Key file on the Endpoint.
Skip SSL Verification.	No	Use SSL but skip chain & host verification. Expected: True/False.

ClickHouse

ClickHouse is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Servers URL	Yes	http://127.0.0.1:8123
User name	No	User name for the ClickHouse service.
Password	No	Password

NTPD

NTPD is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

MySQL

MySQL is supported in vRealize Operations Manager.

To activate the MySQL plug-in and fetch the credentials, refer to the article Steps to fetch password for telegraf user of MySQL, while activating plugin (81153) at the VMware Support Knowledge Base.

Use the port number 3306 to run MySQL and the telegraf credentials and activate the plug-in.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Port	Yes	The port where MySQL is running. Example: 3306
User name	Yes	User name for the MySQL service. Example: Root
password	Yes	Password
SSL CA	No	Path to the SSL CA file on the Endpoint
SSL Certificate	No	Path to the SSL Certificate file on the Endpoint
SSL Key	No	Path to the SSL Key file on the Endpoint.
Hostname	No	Optional hostname for the MySQL Service
Databases	No	Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated. For example, 'database1','database2','database3'.
TLS Connection	No	Accepted values are true, false, and skip-verify.

Redis

Redis is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Redis URL	Yes	servers = ["tcp://localhost:6379"]
SSL CA	No	Secure Socket Layer Certification Authority.
SSL Certificate	No	Secure Socket Layer Certificate.
SSL Key	No	Secure Socket Layer Key
Skip SSL Verification.	No	Skips verification for SSL.

VeloCloud Gateway

VeloCloud Gateway and the following services are supported in vRealize Operations Manager

- Network Time Protocol
- VeloCloud Gateway

In VeloCloud Gateway, we monitor the following processes. For each of these processes, we display a metric which indicates the process status.

- bgpd
- watchquagga
- gwd
- mgd
- natd
- ssh
- vc procmon
- vcsyscmd

VeloCloud Gateway details.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the VeloCloud Gateway instance.

NTPD

NTPD is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.

Pre-Requirements for Application Services

For telegraf agent to collect metrics for some of the application services, you must make modifications in the endpoint VMs. After you make these modifications, the agent will start collecting metrics. You must SSH to the virtual machine where you have deployed the agent and modify the configuration files.

Apache HTTPD

Modify the conf file available in `/etc/httpd/conf.modules.d/status.conf` and enable the mod_status for the HTTPD plugin for the agent to collect metrics.

```
<IfModule mod_status.c>

<Location /server-status>

    SetHandler server-status

</Location>

ExtendedStatus On

</IfModule>
```

If the conf file is not available, you must create one. Restart the HTTPD service after modifying the conf file with the following command:

```
systemctl restart httpd
```

Java Plugins

To monitor Java applications, you can deploy the Jolokia plugin as a .WAR file or .JAR file. If you are deploying a .WAR file, you do not have to restart the services.

For a .JAR file deployment, you have to restart the application service after including the full file path of the JAR in the JMX argument of the JAVA process which you are monitoring.

Nginx

Add the following lines to the conf file available in `/etc/nginx/nginx.conf`:

```
http {
    server {
        location /status {
            stub_status on;
    access_log off;
    allow all;
  }
    }
}
```

Restart the Nginx service with the following command:

```
systemctl restart nginx
```

Postgres

In the configuration file available in the `/var/lib/pgsql/data/pg_hba.conf`, change the value of `local all postgres peer` to `local all postgres md5` and restart the service with the following command:

```
sudo service postgresql restart
```

Cassandra Database

To monitor the Cassandra database application, the Jolokia jar must be included as a JVM input to the Cassandra database application. Complete the following steps:

1   Modify `/etc/default/cassandra`.

```
echo "export JVM_EXTRA_OPTS=\"-javaagent:/usr/share/java/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=localhost\"" | sudo tee -a /etc/default/cassandra
```

2   Alternatively, you can enable the agent by modifying `cassandra-env.sh`. Include the following line at the end of the `cassandra-env.sh`:

```
JVM_OPTS="$JVM_OPTS -javaagent:/usr/share/java/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=localhost"
```

After you see the JVM inputs, restart the Cassandra service.

Oracle Database

To monitor the Oracle database, complete these steps:

1   Download the instant client library from: `https://www.oracle.com/database/technologies/instant-client/downloads.html`.

You must download the Oracle instant library and included it in the PATH.

2   Create a User.

```
CREATE USER <UserName> IDENTIFIED BY <yourpassword>;
 GRANT select_catalog_role TO <UserName>;
 GRANT CREATE SESSION TO <UserName>;
```

3   Install Python 3.6 or later.

```
python3 -m pip install cx_Oracle --upgrade
```

4   Set the PATH of `TNS_ADMIN`.

For example, the path for `TNS_ADMIN` will be similar to `c:\app\product\<version>\dbhome_1\NETWORK\ADMIN"`.

Note Oracle database cannot be activated on Linux platforms.

Active MQ 5.16 and Later Versions

To activate Active MQ 5.16 and later versions, complete the following steps:

- Navigate to `/opt/activemq/apache-activemq-5.16.0/webapps/api/WEB-INF/classes/jolokia-access.xml`

- Remove or comment out the following lines:

  ```
  <cors>
     <strict-checking/>
  </cors>
  ```

- Restart the Active MQ service.

Additional Operations from the Manage Agents Tab

After you have configured vRealize Application Remote Collector and mapped it to a vCenter Server, and installed an agent, you can manage the agents on the VMs from the **Manage Agents** tab. You can view the data centers, hosts, and clusters available in the vCenter Servers you have mapped to vRealize Application Remote Collector. You can start, stop, and update, and uninstall 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 **Administration**, and then from the left pane select **Inventory**. From the right pane, click the **Manage Agents** tab.

Table 1-22. Options

Options	Description
Install	Installs the agents on the selected VM. Select the VMs on which you want to install the agent and click the **Install** icon. For more information, see Install an Agent from the UI.
Uninstall	Uninstalls the agent. Select the VMs on which you want to uninstall the agent and click the **Uninstall** icon. For more information, see Uninstall an Agent.
Update	Updates agents that are at a lower version. Select the VMs on which you want to update the agent and click the **Update** icon. After the agents are updated, the last operation status changes to **Content Upgrade Success**.
Start	If you have temporarily stopped sending metrics to vRealize Operations Manager, 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 vRealize Operations Manager. Select the VMs on which you want to stop the agent and click the **Stop** icon.
Manage Service	You can configure and activate the application services that are discovered on the virtual machines where the agents are installed. For configuration details of each application, see Configuring Supported Application Services.
Manage Service > Remote Check	Allows you to enable remote checks such as ICMP Check, UDP Check, TCP Check, and HTTP Check.

Table 1-22. Options (continued)

Options	Description
Manage Service > Custom Script	Allows you to run custom scripts in the VM and collect custom data which can be then consumed as a metric. For more information, see Custom Script.
Show Detail	Displays the **Summary** tab of the selected VM.
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 1-23. Data Grid Options

Option	Description
VM Name	Name of the virtual machine.
Operating System	Operating system installed on the VM.

Table 1-23. Data Grid Options (continued)

Option	Description
Services Discovered/Configured	List of the supported application services discovered on the VM.
	■ A red dot against the application service indicates that the application service has been activated but there is a problem with data collection.
	When there is more than one application service of the same kind, and one of them is activated, but the other is not collecting data, a red dot is still displayed against the application service.
	■ A gray dot before the application service indicates that the agent requires reactivation. The application service must be reactivated. For reactivation, see Activate an Application Service for more information.
	■ A gray pause symbol indicates that the agents have stopped.
	■ A green icon against the application service indicates that the application service is activated.
	You see a blue icon with three horizontal dots if there is a problem with the activation. Click the question mark for more information about the warning. The warning is also displayed in the following locations:
	■ In the **Objects** tab for the specific application service. Move your cursor over the green icon in the **Collection Status** column.
	■ For the specific application service, click the **Show Details** option from the **Manage Agents** tab. Move your cursor over the green icon in the top panel to view the warning message.
	■ If an application service has been deactivated or not activated, you see a gray pause 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.
	Click the colored dots for more information about the application services.
Agent Status	Displays the status of the agent at the end point.
	■ Blue icon. Indicates that the agent is not installed.
	■ Green icon. Indicates that the agent is running.
	■ Red icon. Indicates that the agent has stopped.
	■ Gray dot. Appears in front of the service and indicates that plugin reactivation is required.

Table 1-23. Data Grid Options (continued)

Option	Description
Last Operation Status	Status of the last operation. The possible values are: ■ 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 ■ Download Success
VM State	Power status of the VMs. The possible values are: ■ Powered On ■ Powered Off
ARC	FQDN of the instance of the vRealize Application Remote Collector that you are using.
Agent Version	Version of the 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.

To manage the agent, follow these steps:

1 Install the agent.

For more information, see Install an Agent from the UI.

2 Manage the application services on each agent.

For more information, see Configure 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.

Note   You cannot run the vRealize Application Remote Collector agent on the same VM as the End Point Operations Management agent.

Custom Script

You can run custom scripts in the VM and collect custom data which can then be consumed as a metric.

Prerequisites

■   All the scripts that you run using the custom script, must output a single integer value. If the output is not a single integer value, an error is displayed in the user interface.

■   The custom script uses Telegraf's `exec` plugin to run scripts on a VM's operating system. The scripts are run by the user who installed the Telegraf agent on an operating system. In Linux operating systems, a special user called *arcuser* with specific privileges, is created for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that *arcuser* user. Ensure that the *arcuser* can run the scripts that use the custom script (the *arcuser* must have permissions to run the script). For example, the *arcuser* created automatically by vRealize Application Remote Collector, does not have privileges to run scripts which are stored under the `/root` directory.

■   The script must be placed in the `/opt/vmware` folder.

Instance Settings

Option	Description
Status	Enable the custom script execution.
Display Name	Add a suitable name for the script. The * is an invalid character and must not be used in the name.
Filepath	Enter the path to the script file on the end point VM.
Prefix	Enter a prefix if necessary.
Args	List the arguments in the script.
Timeout	Enter a script execution timeout on the VM.

After you save the script, it appears in the left pane of the **Custom Script** dialog box. You can add or delete scripts by clicking the **Add** or **Delete** buttons in the left pane. After the scripts have been added and saved, from the **Manage Agents tab > Services Discovered/Configured** column, you see the **Custom Script** label. Point to the **Custom Script** label to view the list of scripts and their status.

Note

- The custom script must throw all errors in the format `ERROR|<Error_message>` for the error propagation to work. If the script does not throw an error in the given format, vRealize Operations Manager displays an error message `Unable to parse the error message. Please check the endpoint` in the user interface. This is by design, until vRealize Application Remote Collector propagates the exact error message.

- The bash script must start with `shebang (#!/bin/bash)`.

All Metrics Tab

When data is collected successfully, you can view the script as a metric for the VM, in the **All Metrics** tab. The script metrics are created under an object called `Custom Script` which is a single object per VM. All the metrics from the scripts for the VM are placed under that `Custom Script` object that contains all the custom scripts you have created. You can view the output for the specific metric. The metric name under the `Scripts` folder is the display name that the user specifies while creating the script configuration. For example, if you set the display name as **Python script**, then a metric is created with the name **Python script** if data is collected successfully.

Deactivate an Application Service

You can deactivate an application service to stop monitoring the application service that is sending data to vRealize Operations Manager.

Prerequisite

- If plugin deactivation requires the location of a file (for example, client certificates for SSL Trust) on the endpoint VM, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

  Note If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

Deactivate an Application Service

To deactivate a plugin to stop monitoring the application service that is sending data to vRealize Operations Manager, complete the following steps:

1   Navigate to the **Inventory > Manage Agents** tab.

2   Select the VM on which the agent is already installed.

3   Select the **Manage Service** icon and then from the drop-down menu select the **service name**.

4   Deactivate the application service from the right pane of the **Manage <service name> Agent** dialog box.

5   Click **Save**.

When you stop an agent, you cannot activate or deactivate a plugin. If the VM is powered off or if you lose connection with vRealize Application Remote Collector, you cannot configure or activate a plugin.

For information on activating an application service, see Activate an Application Service.

Uninstall an Agent

You must select the VMs on which you want to uninstall the agent.

Prerequisites

- Time synchronization between vRealize Application Remote Collector, vRealize Operations Manager, ESX hosts, and Windows and Linux target VMs is mandatory for secure communication.

- Guest operation privileges are required to install agents on virtual machines. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.

- Account privilege prerequisites. See User Account Prerequisites for more details.

- End-point VM configuration requirements.

  - Linux requirements

    Commands: `/bin/bash, sudo, tar, awk, curl`

    Packages: `coreutils (chmod, chown, cat), shadow-utils (useradd, groupadd, userdel, groupdel)`

    Configure mount point on `/tmp` directory to allow script execution.

  - Windows 2012 R2 requirement

    The end point must be updated with the Universal C Runtime. Refer to the following link for more information.

  - Windows requirement

    The Visual C++ version must be higher than 14.

- VMware Tools must be installed and running on the VM on which you want to install the agent.

Procedure

1   From the **Manage Agents** tab, click the **Uninstall** icon. 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 **Uninstall 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.

UAC Enabled Machines on Windows End Points

The bits are downloaded to the end point. You have to manually uninstall the bits.

    a   From `C:\VMware\UCP\downloads`, run a bootstrap launcher.

    b   Go to `%SYSTEMDRIVE%\VMware\UCP\downloads`.

    c   Open `cmd` with administrator privileges.

    d   Run the `cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1` command.

    e   View the results from `uaf_bootstrap.log`.

    f   Verify the status of agent uninstallation from the **Agent Status** and **Last Operation Status** columns in the **Manage Agents**

For information about installing an agent, see Install an Agent from the UI.

Configure Application Services

You can configure the application services 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** tab.

2   Select **Manage Service** and then from the drop-down menu select the **service name**. You see the **Manage <service name> Agent** 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. For configuration details of each application, see Configuring Supported Application Services.

6   Click **Save** and then **Close**.

Fields with a star are mandatory.

For more information about the status details that appear against the application services in the **Services Discovered/Configured** column, see the table called Data Grid Options in Additional Operations from the Manage Agents Tab.

What to do next

You can monitor the applications services from vRealize Operations Manager.

[Optional] Backup and Restore a vRealize Application Remote Collector Instance

You can run the backup and restore script to ensure that VMware vRealize Operations Manager continues to receive data after the vRealize Application Remote Collector instance becomes unavailable. All the existing endpoints that are configured will automatically connect back to vRealize Application Remote Collector and continue to send data after you restore the vRealize Application Remote Collector instance. This is an optional task which you can run if you are facing issues with the vRealize Application Remote Collector appliance.

The task is divided into two parts. The first part involves performing an on-demand back up of the vRealize Application Remote Collector connection and configuration details. A cron job also performs the back up automatically every day.

The second part involves restoring the vRealize Application Remote Collector instance using the backup file that you created, or the backup file created by the cron job.

Prerequisites

- vRealize Application Remote Collector appliance must be configured with a static I.P. or static FQDN. The endpoints must be configured.

- Back up the network configuration details of the vRealize Application Remote Collector appliance. Capture the network configuration details of vRealize Application Remote Collector either using the VAMI UI or vCenter Server Tools. Keep the network details available when you restore the vRealize Application Remote Collector appliance from the backup.

- The sizing of the new vRealize Application Remote Collector appliance that you are restoring a backup to, should be greater or equal to the old appliance. The network configuration, static I.P. or static FQDN should be the same. This is to enable the endpoint VMs to reach the new appliance.

Procedure

1 Back up a running instance of vRealize Application Remote Collector by making a copy of the connection and configuration details.

  a Connect to the virtual machine running vRealize Application Remote Collector using SSH.

  b Enter the following command to access the scripts folder:

```
cd /ucp/ucp-config-scripts
```

  c Run the `arc-state-bundle.sh` script with the backup option. The script performs a back up or restore task based on the option you provide.

```
./arc-state-bundle.sh backup_state
```

   Running this script pushes the backup file to the `/ucp-bkup/state-bundles` folder. The filename is in the format `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar`. This file contains the connection and configuration details for the endpoints.

  d Archive the `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file to a remote location.

2 A cron job also runs every day and backs up the `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file. The `.tar` file is stored for five days. On the sixth day, the oldest .tar file is deleted and replaced. In order to restore the vRealize Application Remote Collector appliance from the .tar file, archive the file to a remote location.

3 Restore the backed up configuration files to a new vRealize Application Remote Collector appliance.

  a Configure the new vRealize Application Remote Collector appliance with the same network and IP configuration as the previous appliance. This information is available in the network configuration file that you backed up.

  b Connect to the VM running vRealize Application Remote Collector using SSH.

  c Retrieve the latest `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file from the archive, and copy it to a location which is accessible by the vRealize Application Remote Collector appliance.

d   Enter the following command to access the scripts folder:

```
cd /ucp/ucp-config-scripts
```

e   Run the `arc-state-bundle.sh` script. Use the restore option. Provide the location of the `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file.

```
./arc-state-bundle.sh restore_state <<location of the backed up tar file, with the
filename.tar extension>>
```

The above command looks for the file starting with `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` to load. The script configures the new vRealize Application Remote Collector appliance with the same settings as the instance that went down, and restarts all the containers.

For example, the following command restores the appliance from the state bundle `/tmp/fromArchive/Application-Remote-Collector-State-Bundle_2019-04-02-18:31:36.tar` from the `/tmp/fromArchive/` location:

```
./arc-state-bundle.sh restore_state "/tmp/fromArchive/Application-Remote-Collector-
State-Bundle_2019-04-02-18:31:36.tar"
```

Results

The restoration of the vRealize Application Remote Collector is complete, and it is available again. The existing endpoints connect back to vRealize Application Remote Collector and continue to send data.

What to do next

If the vRealize Application Remote Collector instance was sending data to VMware vRealize Operations Manager, then adapter collection might fail when the vRealize Application Remote Collector instance stops working. In the VMware vRealize Operations Manager, the status of the adapter instances changes to indicate that it has failed. If this happens, you must manually start the adapter instance after restoring the vRealize Application Remote Collector appliance.

Summary of Discovered and Supported Operating Systems and Application Services

You can monitor application services and operating systems from vRealize Operations Manager to view services and processes.

Where You View Applications in vRealize Operations Manager

From the menu, select **Home**, and then in the left pane select **Monitor Applications**.

Discovered Operating Systems and Services

You see the application services that are discovered on the virtual machines where the agents are installed. From the **Discovered Operating Systems and Services** section in the **Monitor Applications** page, click the text next to the number to view the status of the agent, the operation status, the power status of the VM, and the list of supported application services discovered on the VM. For more information, see Additional Operations from the Manage Agents Tab.

Supported Operating Systems

You see a list of supported operating systems for which vRealize Operations Manager collects metrics.

Supported Services

You see a list of supported services for which vRealize Operations Manager collects metrics.

Metrics Collected

Metrics are collected for operating systems, application services, and remote checks.

Operating System Metrics

Metrics are collected for Linux and Windows operating systems. The metrics are collected after the vRealize Application Remote Collector agent is deployed on the VM.

Linux Platforms

The following metrics are collected for Linux operating systems:

Table 1-24. Metrics for Linux

Metric	Metric Category	KPI
<Instance name>| Usage Idle	CPU	False
<Instance name>| Usage IO-Wait	CPU	False
<Instance name>|Time Active	CPU	True
<Instance name>|Time Guest	CPU	False
<Instance name>|Time Guest Nice	CPU	False
<Instance name>|Time Idle	CPU	False
<Instance name>|Time IO-Wait	CPU	False
<Instance name>|Time IRQ	CPU	True
<Instance name>|Time Nice	CPU	False
<Instance name>|Time Soft IRQ	CPU	True
<Instance name>|Time Steal	CPU	False

Table 1-24. Metrics for Linux (continued)

Metric	Metric Category	KPI
<Instance name>|Time System	CPU	False
<Instance name>|Time User	CPU	True
<Instance name>|Usage Active (%)	CPU	True
<Instance name>|Usage Guest (%)	CPU	False
<Instance name>|Usage Guest Nice (%)	CPU	False
<Instance name>|Usage IRQ (%)	CPU	True
<Instance name>|Usage Nice (%)	CPU	False
<Instance name>|Usage Soft IRQ (%)	CPU	True
<Instance name>|Usage Steal (%)	CPU	False
<Instance name>|Usage System (%)	CPU	True
<Instance name>|Usage User (%)	CPU	True
IO Time	Disk	False
Read Time	Disk	False
Reads	Disk	False
Write Time	Disk	False
Writes	Disk	False
<Instance name>|Disk Free	Disk	False
<Instance name>|Disk Total	Disk	False
<Instance name>|Disk Used (%)	Disk	False
Cached	Memory	False
Free	Memory	False
Inactive	Memory	False
Total	Memory	True
Used	Memory	True
Used Percent	Memory	True
Blocked	Processes	True
Dead	Processes	False
Running	Processes	False

Table 1-24. Metrics for Linux (continued)

Metric	Metric Category	KPI
Sleeping	Processes	False
Stopped	Processes	False
Zombies	Processes	False
Free	Swap	False
In	Swap	False
Out	Swap	False
Total	Swap	True
Used	Swap	True
Used Percent	Swap	True

Windows Platforms

The following metrics are collected for Windows operating systems:

Table 1-25. Metrics for Windows

Metric	Metric Category	KPI
Idle Time	CPU	False
Interrupt Time	CPU	False
Interrupts persec	CPU	True
Privileged Time	CPU	False
Processor Time	CPU	False
User Time	CPU	False
Avg. Disk Bytes Read	Disk	False
Avg. Disk sec Read	Disk	False
Avg. Disk sec Write	Disk	False
Avg. Disk Write Queue Length	Disk	False
Avg. Disk Read Queue Length	Disk	False
Disk Read Time	Disk	False
Disk Write Time	Disk	False
Free Megabytes	Disk	False
Free Space	Disk	False

Table 1-25. Metrics for Windows (continued)

Metric	Metric Category	KPI
Idle Time	Disk	False
Split IO persec	Disk	False
Available Bytes	Memory	True
Cache Bytes	Memory	False
Cache Faults persec	Memory	False
Committed Bytes	Memory	True
Demand Zero Faults persec	Memory	False
Page Faults persec	Memory	True
Pages persec	Memory	False
Pool Nonpaged Bytes	Memory	True
Pool Paged Bytes	Memory	False
Transition Faults persec	Memory	False
Elapsed Time	Process	False
Handle Count	Process	False
IO Read Bytes persec	Process	False
IO Read Operations persec	Process	False
IO Write Bytes persec	Process	False
IO Write Operations persec	Process	False
Privileged Time	Process	False
Processor Time	Process	False
Thread Count	Process	False
User Time	Process	False
Context Switches persec	System	False
Processes	System	False
Processor Queue Length	System	False
System Calls persec	System	False
System Up Time	System	False
Threads	System	False

Application Service Metrics

Metrics are collected for 20 application services.

Active Directory Metrics
Metrics are collected for the Active Directory application service.

Table 1-26. Active Directory Metrics

Metric Name	Category	KPI
Database Cache % Hit (%)	Active Directory Database	True
Database Cache Page Faults/sec	Active Directory Database	True
Database Cache Size	Active Directory Database	False
Data Lookups	Active Directory DFS Replication	False
Database Commits	Active Directory DFS Replication	True
Avg Response Time	Active Directory DFSN	True
Requests Failed	Active Directory DFSN	False
Requests Processed	Active Directory DFSN	False
Dynamic Update Received	Active Directory DNS	False
Dynamic Update Rejected	Active Directory DNS	False
Recursive Queries	Active Directory DNS	False
Recursive Queries Failure	Active Directory DNS	False
Secure Update Failure	Active Directory DNS	False
Total Query Received	Active Directory DNS	True
Total Response Sent	Active Directory DNS	True
Digest Authentications	Active Directory Security System-Wide Statistics	True
Kerberos Authentications	Active Directory Security System-Wide Statistics	True
NTLM Authentications	Active Directory Security System-Wide Statistics	True
Directory Services:<InstanceName>|Base Searches persec	Active Directory Services	False
Directory Services:<InstanceName>|Database adds persec	Active Directory Services	False
Directory Services:<InstanceName>|Database deletes persec	Active Directory Services	False
Directory Services<InstanceName>|Database modifys/sec	Active Directory Services	False

Table 1-26. Active Directory Metrics (continued)

Metric Name	Category	KPI
Directory Services<InstanceName>|Database recycles/sec	Active Directory Services	False
Directory Services<InstanceName>|DRA Inbound Bytes Total/sec	Active Directory Services	False
Directory Services<InstanceName>|DRA Inbound Objects/sec	Active Directory Services	False
Directory Services<InstanceName>|DRA Outbound Bytes Total/sec	Active Directory Services	False
Directory Services<InstanceName>|DRA Outbound Objects/sec	Active Directory Services	False
Directory Services<InstanceName>|DRA Pending Replication Operations	Active Directory Services	False
Directory Services<InstanceName>|DRA Pending Replication Synchronizations	Active Directory Services	False
Directory Services<InstanceName>|DRA Sync Requests Made	Active Directory Services	False
Directory Services<InstanceName>|DRA Sync Requests Successful	Active Directory Services	False
Directory Services<InstanceName>|DS Client Binds/sec	Active Directory Services	True
Directory Services<InstanceName>|DS Directory Reads/sec	Active Directory Services	False
Directory Services<InstanceName>|DS Directory Searches/sec	Active Directory Services	True
Directory Services<InstanceName>|DS Server Binds/sec	Active Directory Services	True
Directory Services<InstanceName>|DS Threads in Use	Active Directory Services	True
Directory Services:<InstanceName>|LDAP Active Threads	Active Directory Services	False
Directory Services:<InstanceName>|LDAP Client Sessions	Active Directory Services	True
Directory Services<InstanceName>|LDAP Closed Connections/sec	Active Directory Services	False
Directory Services<InstanceName>|LDAP New Connections/sec	Active Directory Services	True
Directory Services<InstanceName>|LDAP Searches/sec	Active Directory Services	True
Directory Services<InstanceName>|LDAP Successful Binds/sec	Active Directory Services	False

Table 1-26. Active Directory Metrics (continued)

Metric Name	Category	KPI
Directory Services<InstanceName>| LDAP UDP operations/sec	Active Directory Services	False
Directory Services:<InstanceName>| LDAP Writes/sec	Active Directory Services	False

No metrics are collected for the category Active Directory.

Apache Tomcat

Metrics are collected for the Apache Tomcat application service.

Table 1-27. Apache Tomcat

Metric Name	Category	KPI
Buffer Pool<InstanceName>|Count	Tomcat Server	False
Buffer Pool<InstanceName>|Memory Used	Tomcat Server	False
Buffer Pool<InstanceName>|Total Capacity	Tomcat Server	False
Class Loading|Loaded Class Count	Tomcat Server	False
Class Loading|Total Loaded Class Count	Tomcat Server	False
Class Loading|Unloaded Class Count	Tomcat Server	False
File Descriptor Usage|Max File Descriptor Count	Tomcat Server	False
File Descriptor Usage|Open File Descriptor Count	Tomcat Server	False
Garbage Collection:<InstanceName>| Total Collection Count	Tomcat Server	False
Garbage Collection:<InstanceName>| Total Collection Time	Tomcat Server	True
JVM Memory|Heap Memory Usage| Committed Memory	Tomcat Server	False
JVM Memory|Heap Memory Usage| Initial Memory	Tomcat Server	False
JVM Memory|Heap Memory Usage| Maximum Memory	Tomcat Server	False
JVM Memory|Heap Memory Usage| Used Memory	Tomcat Server	False
JVM Memory|Non Heap Memory Usage|Committed Memory	Tomcat Server	False
JVM Memory|Non Heap Memory Usage|Initial Memory	Tomcat Server	False

Table 1-27. Apache Tomcat (continued)

Metric Name	Category	KPI
JVM Memory|Non Heap Memory Usage|Maximum Memory	Tomcat Server	False
JVM Memory|Non Heap Memory Usage|Used Memory	Tomcat Server	False
JVM Memory|Number of Object Pending Finalization Count	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Peak Usage|Committed Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Peak Usage|Initial Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Peak Usage|Maximum Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Peak Usage|Used Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Usage|Committed Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Usage|Initial Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Usage|Maximum Memory	Tomcat Server	False
JVM Memory|Pool:<InstanceName>|Usage|Used Memory	Tomcat Server	False
Process CPU Usage (%)	Tomcat Server	True
System CPU Usage (%)	Tomcat Server	True
System Load Average (%)	Tomcat Server	True
Threading|Thread Count	Tomcat Server	False
Uptime	Tomcat Server	True
JSP Count	Tomcat Server Web Module	False
JSP Reload Count	Tomcat Server Web Module	False
JSP Unload Count	Tomcat Server Web Module	False
Servlet:<InstanceName>|Total Request Count	Tomcat Server Web Module	False
Servlet:<InstanceName>|Total Request Error Count	Tomcat Server Web Module	False
Servlet:<InstanceName>|Total Request Processing Time	Tomcat Server Web Module	False
Cache : Hit Count	Tomcat Server Web Module	False

Table 1-27. Apache Tomcat (continued)

Metric Name	Category	KPI
Cache : Lookup Count	Tomcat Server Web Module	False
Current Thread Count	Tomcat Server Global Request Processor	True
Current Threads Busy	Tomcat Server Global Request Processor	True
errorRate	Tomcat Server Global Request Processor	False
Total Request Bytes Received	Tomcat Server Global Request Processor	False
Total Request Bytes Sent	Tomcat Server Global Request Processor	False
Total Request Count	Tomcat Server Global Request Processor	True
Total Request Error Count	Tomcat Server Global Request Processor	True
Total Request Processing Time	Tomcat Server Global Request Processor	False

MS SQL Metrics
Metrics are collected for the MS SQL application service.

Table 1-28. MS SQL Metrics

Metric Name	Category	KPI
CPU<InstanceName>|CPU Usage (%)	Microsoft SQL Server	False
Database IO|Rows Reads Bytes/Sec	Microsoft SQL Server	False
Database IO|Rows Reads/Sec	Microsoft SQL Server	False
Database IO|Rows Writes Bytes/Sec	Microsoft SQL Server	False
Database IO|Rows Writes/Sec	Microsoft SQL Server	False
Performance|Access Methods|Full Scans per second	Microsoft SQL Server	False
Performance|Access Methods|Index Searches	Microsoft SQL Server	False
Performance|Access Methods|Page Splits per second	Microsoft SQL Server	False
Performance|Broker Activation|Stored Procedures Invoked per second	Microsoft SQL Server	False
Performance|Buffer Manager|Buffer cache hit ratio (%)	Microsoft SQL Server	True

Table 1-28. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance|Buffer Manager|Checkpoint Pages/sec	Microsoft SQL Server	True
Performance|Buffer Manager|Lazy writes per second	Microsoft SQL Server	True
Performance|Buffer Manager|Page life expectancy	Microsoft SQL Server	True
Performance|Buffer Manager|Page lookups per second	Microsoft SQL Server	False
Performance|Buffer Manager|Page reads per second	Microsoft SQL Server	False
Performance|Buffer Manager|Page writes per second	Microsoft SQL Server	False
Performance|Databases|Active Transactions	Microsoft SQL Server	True
Performance|Databases|Data File(s) Size	Microsoft SQL Server	True
Performance|Databases|Log Bytes Flushed/Sec	Microsoft SQL Server	False
Performance|Databases|Log File(s) Size	Microsoft SQL Server	False
Performance|Databases|Log File(s) Used Size	Microsoft SQL Server	False
Performance|Databases|Log Flush Wait Time	Microsoft SQL Server	False
Performance|Databases|Log Flushes per second	Microsoft SQL Server	False
Performance|Databases|Transactions per second	Microsoft SQL Server	False
Performance|Databases|Write Transactions per second	Microsoft SQL Server	False
Performance|Databases|XTP Memory Used	Microsoft SQL Server	False
Performance|General Statistics|Active temp Tables	Microsoft SQL Server	False
Performance|General Statistics|Logins per second	Microsoft SQL Server	False
Performance|General Statistics|Logouts per second	Microsoft SQL Server	False
Performance|General Statistics|Processes Blocked	Microsoft SQL Server	False

Table 1-28. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance|General Statistics|Temp Tables Creation Rate	Microsoft SQL Server	False
Performance|General Statistics|User Connections	Microsoft SQL Server	False
Performance|Locks|Average Wait Time	Microsoft SQL Server	False
Performance|Locks|Lock Requests per second	Microsoft SQL Server	False
Performance|Locks|Lock Wait Time	Microsoft SQL Server	True
Performance|Locks|Lock Waits per second	Microsoft SQL Server	True
Performance|Locks|Number of Deadlocks per second	Microsoft SQL Server	True
Performance|Memory Manager|Connection Memory	Microsoft SQL Server	False
Performance|Memory Manager|Lock Memory	Microsoft SQL Server	False
Performance|Memory Manager|Log Pool Memory	Microsoft SQL Server	False
Performance|Memory Manager|Memory Grants Pending	Microsoft SQL Server	True
Performance|Memory Manager|SQL Cache Memory	Microsoft SQL Server	False
Performance|Memory Manager|Target Server Memory	Microsoft SQL Server	True
Performance|Memory Manager|Total Server Memory	Microsoft SQL Server	True
Performance|Resource Pool Stats|internal|Active memory grant amount	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|CPU Usage Percentage (%)	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|Disk Read Bytes per second	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|Disk Read IO	Microsoft SQL Server	False
Wait Stats:<InstanceName>|Wait Time (ms)	Microsoft SQL Server	False
Wait Stats<InstanceName>|Number of Waiting tasks (ms)	Microsoft SQL Server	False

Table 1-28. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance|Resource Pool Stats|internal|Disk Read IO Throttled Per Second	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|Disk Write Bytes per second (Bps)	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|Disk Write IO Throttled per second	Microsoft SQL Server	False
Performance|Resource Pool Stats|internal|Used Memory	Microsoft SQL Server	False
Performance|SQL Statistics | Batch Requests Per Second	Microsoft SQL Server	False
Performance|SQL Statistics | SQL Compilations per second	Microsoft SQL Server	False
Performance|SQL Statistics | SQL Re-Compilations per second	Microsoft SQL Server	False
Performance|Transactions | Free space in tempdb (KB)	Microsoft SQL Server	False
Performance|Transactions | Transactions	Microsoft SQL Server	False
Performance|Transactions | Version Store Size (KB)	Microsoft SQL Server	False
Performance|User Settable Counter | User Counter 0 to 10	Microsoft SQL Server	False
Performance|Workload Group Stats|internal|Active Requests	Microsoft SQL Server	False
Performance|Workload Group Stats|internal|Blocked Tasks	Microsoft SQL Server	False
Performance|Workload Group Stats|internal|CpU Usage (%)	Microsoft SQL Server	False
Performance|Workload Group Stats|internal|Queued Requests	Microsoft SQL Server	False
Performance|Workload Group Stats|internal|Request Completed/sec	Microsoft SQL Server	False

There are no metrics collected for Microsoft SQL Server Database.

PostgresSQL

Metrics are collected for the PostgresSQL application service.

Table 1-29. PostgresSQL

Metric Name	Category	KPI
Buffers|Buffers Allocated	PostgresSQL	False
Buffers|Buffers Written by Backend	PostgresSQL	True
Buffers|Buffers Written by Background Writer	PostgresSQL	True
Buffers|Buffers Written During Checkpoints	PostgresSQL	True
Buffers|fsync Call Executed by Backend	PostgresSQL	False
Checkpoints|Checkpoints sync time	PostgresSQL	False
Checkpoints|Checkpoints write time	PostgresSQL	False
Checkpoints|Requested checkpoints performed count	PostgresSQL	False
Checkpoints|Scheduled checkpoints performed count	PostgresSQL	False
Clean scan stopped count	PostgresSQL	False
Disk Blocks|Blocks Cache Hits	PostgresSQL Database	False
Disk Blocks|Blocks Read	PostgresSQL Database	False
Disk Blocks|Blocks Read Time	PostgresSQL Database	False
Disk Blocks|Blocks Write Time	PostgresSQL Database	False
Statistics|Backends Connected	PostgresSQL Database	False
Statistics|Data Written by Queries	PostgresSQL Database	True
Statistics|Deadlocks Detected	PostgresSQL Database	True
Statistics|Queries Cancelled	PostgresSQL Database	True
Statistics|Temp Files Created by Queries	PostgresSQL Database	False
Transactions|Transactions Committed	PostgresSQL Database	True
Transactions|Transactions Rolled Back	PostgresSQL Database	True
Tuples|Tuples Deleted	PostgresSQL Database	True
Tuples|Tuples Fetched	PostgresSQL Database	True
Tuples|Tuples Inserted	PostgresSQL Database	True

Table 1-29. PostgresSQL (continued)

Metric Name	Category	KPI
Tuples|Tuples Returned	PostgresSQL Database	True
Tuples|Tuples Updated	PostgresSQL Database	True

IIS Metrics

Metrics are collected for the IIS application service.

Table 1-30. IIS Metrics

Metric Name	Category	KPI
HTTP Service Request Queues<InstanceName>AppPool|CurrentQueueSize	IIS HTTP Service Request Queues	True
HTTP Service Request Queues<InstanceName>AppPool|RejectedRequests	IIS HTTP Service Request Queues	False
Web Services<InstanceName> Web Site|Bytes Received	IIS Web Services	False
Web Services<InstanceName> Web Site|Bytes Sent/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Bytes Total/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Connection Attempts/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Current Connections	IIS Web Services	False
Web Services<InstanceName> Web Site|Get Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Locked Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Not Found Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Post Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site|Service Uptime	IIS Web Services	False
Web Services<InstanceName> Web Site|Total Bytes Sent	IIS Web Services	False
Web Services<InstanceName> Web Site|Total Get Requests	IIS Web Services	True
Web Services<InstanceName> Web Site|Total Post Requests	IIS Web Services	True

Table 1-30. IIS Metrics (continued)

Metric Name	Category	KPI
Web Services<InstanceName> Web Site|Total Put Requests	IIS Web Services	False
Current File Cache Memory Usage (bytes)	IIS Web Services Cache	False
File Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Misses	IIS Web Services Cache	False
Total Flushed URIs	IIS Web Services Cache	False
URI Cache Hits	IIS Web Services Cache	False
URI Cache Hits Percent (%)	IIS Web Services Cache	False
URI Cache Misses	IIS Web Services Cache	False
ASP.NET<InstanceName>|Application Restarts	IIS ASP.NET	True
ASP.NET<InstanceName>|Request Wait Time	IIS ASP.NET	True
ASP.NET<InstanceName>|Requests Current	IIS ASP.NET	True
ASP.NET<InstanceName>|Requests Queued	IIS ASP.NET	True
ASP.NET<InstanceName>|Requests Rejected	IIS ASP.NET	True
MS.NET<InstanceName>|Allocated Bytes/sec	MS.NET	True
MS.NET<InstanceName>|Current Queue Length	MS.NET	False
MS.NET<InstanceName>|Finalization Survivors	MS.NET	False
MS.NET<InstanceName>|Gen 0 Collections	MS.NET	False
MS.NET<InstanceName>|Gen 0 heap size	MS.NET	False
MS.NET<InstanceName>|Gen 1 Collections	MS.NET	False
MS.NET<InstanceName>|Gen 1 heap size	MS.NET	False
MS.NET<InstanceName>|Gen 2 Collections	MS.NET	False

Table 1-30. IIS Metrics (continued)

Metric Name	Category	KPI
MS.NET<InstanceName>|Gen 2 heap size	MS.NET	False
MS.NET<InstanceName>|IL Bytes Jitted / sec	MS.NET	False
MS.NET<InstanceName>|Induced GC	MS.NET	False
MS.NET<InstanceName>|Large Object Heap size	MS.NET	False
MS.NET<InstanceName>|No of current logical Threads	MS.NET	True
MS.NET<InstanceName>|No of current physical Threads	MS.NET	True
MS.NET<InstanceName>|No of current recognized threads	MS.NET	False
MS.NET<InstanceName>|No of Exceps Thrown / sec	MS.NET	True
MS.NET<InstanceName>|No of total recognized threads	MS.NET	False
MS.NET<InstanceName>|Percent Time in Jit	MS.NET	False
MS.NET<InstanceName>|Pinned Objects	MS.NET	False
MS.NET<InstanceName>|Stack Walk Depth	MS.NET	False
MS.NET<InstanceName>|Time in RT checks	MS.NET	False
MS.NET<InstanceName>|Time Loading	MS.NET	True
MS.NET<InstanceName>|Total No of Contentions	MS.NET	False
MS.NET<InstanceName>|Total Runtime Checks	MS.NET	True

MS Exchange Server Metrics

Metrics are collected for the MS Exchange Server application service.

Table 1-31. MS Exchange Server Metrics

Metric Name	Category	KPI
Active Manager Server|Active Manager Role	MS Exchange	False
Active Manager Server|Database State Info Writes per second	MS Exchange	False

Table 1-31. MS Exchange Server Metrics (continued)

Metric Name	Category	KPI
Active Manager Server|GetServerForDatabase Server-Side Calls	MS Exchange	False
Active Manager Server|Server-Side Calls per second	MS Exchange	True
Active Manager Server|Total Number of Databases	MS Exchange	True
ActiveSync|Average Request Time	MS Exchange	True
ActiveSync|Current Requests	MS Exchange	False
ActiveSync|Mailbox Search Total	MS Exchange	False
ActiveSync|Ping Commands Pending	MS Exchange	False
ActiveSync|Requests per second	MS Exchange	True
ActiveSync|Sync Commands per second	MS Exchange	True
ASP.NET|Application Restarts	MS Exchange	False
ASP.NET|Request Wait Time	MS Exchange	True
ASP.NET|Worker Process Restarts	MS Exchange	False
Autodiscover Service|Requests per second	MS Exchange	True
Availability Service|Average Time to Process a Free Busy Request	MS Exchange	True
Outlook Web Access|Average Search Time	MS Exchange	True
Outlook Web Access|Requests per second	MS Exchange	False
Outlook Web Access|Current Unique Users	MS Exchange	False
Performance|Database Cache Hit (%)	MS Exchange Database	False
Performance|Database Page Fault Stalls per second	MS Exchange Database	True
Performance|I/O Database Reads Average Latency	MS Exchange Database	True
Performance|I/O Database Writes Average Latency	MS Exchange Database	True
Performance|I/O Log Reads Average Latency	MS Exchange Database	False

Table 1-31. MS Exchange Server Metrics (continued)

Metric Name	Category	KPI
Performance|I/O Log Writes Average Latency	MS Exchange Database	False
Performance|Log Record Stalls per second	MS Exchange Database	False
Performance|Log Threads Waiting	MS Exchange Database	False
Performance|I/O Database Reads Average Latency	MS Exchange Database Instance	False
Performance|I/O Database Writes Average Latency	MS Exchange Database Instance	False
Performance|Log Record Stalls per second	MS Exchange Database Instance	False
Performance|Log Threads Waiting	MS Exchange Database Instance	False
Performance|LDAP Read Time	MS Exchange Domain Controller	False
Performance|LDAP Search Time	MS Exchange Domain Controller	False
Performance|LDAP Searches Timed Out per minute	MS Exchange Domain Controller	False
Performance|Long Running LDAP Operations per minute	MS Exchange Domain Controller	False
Performance|Connection Attempts per second	MS Exchange Web Server	True
Performance|Current Connections	MS Exchange Web Server	False
Performance|Other Request Methods per second	MS Exchange Web Server	False
Process|Handle Count	MS Exchange Windows Service	False
Process|Memory Allocated	MS Exchange Windows Service	False
Process|Processor Time (%)	MS Exchange Windows Service	True
Process|Thread Count	MS Exchange Windows Service	False
Process|Virtual Memory Used	MS Exchange Windows Service	False
Process|Working Set	MS Exchange Windows Service	False

JBoss EAP Metrics

Metrics are collected for the JBoss EAP application service.

Table 1-32. JBoss EAP Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName>|Count	Jboss Server	False
Buffer Pool<InstanceName>|Memory Used	Jboss Server	False
Buffer Pool<InstanceName>|Total Capacity	Jboss Server	False
Class Loading|Loaded Class Count	Jboss Server	False
Class Loading|Total Loaded Class Count	Jboss Server	False
Class Loading|Unloaded Class Count	Jboss Server	False
File Descriptor Usage|Max File Descriptor Count	Jboss Server	False
File Descriptor Usage|Open File Descriptor Count	Jboss Server	False
Http Listener<InstanceName>|Bytes Received	Jboss Server	False
Http Listener<InstanceName>|Bytes Sent	Jboss Server	False
Http Listener<InstanceName>|Error Count	Jboss Server	False
Http Listener<InstanceName>| Request Count	Jboss Server	False
Https Listener<InstanceName>|Bytes Received	Jboss Server	False
Https Listener<InstanceName>|Bytes Sent	Jboss Server	False
Https Listener<InstanceName>|Error Count	Jboss Server	False
Https Listener<InstanceName>| Request Count	Jboss Server	False
Process CPU Usage (%)	Jboss Server	False
System CPU Usage (%)	Jboss Server	False
System Load Average (%)	Jboss Server	False
Threading|Daemon Thread Count	Jboss Server	False
Threading|Peak Thread Count	Jboss Server	False
Threading|Thread Count	Jboss Server	False
Threading|Total Started Thread Count	Jboss Server	False

Table 1-32. JBoss EAP Metrics (continued)

Metric Name	Category	KPI
Uptime	Jboss Server	False
UTILIZATION|Heap Memory Usage	Jboss Server	False
Garbage Collection<InstanceName>|Total Collection Count	Jboss JVM Garbage Collector	False
Garbage Collection<InstanceName>|Total Collection Time	Jboss JVM Garbage Collector	False
JVM Memory|Heap Memory Usage|Committed Memory	Jboss JVM Memory	False
JVM Memory|Heap Memory Usage|Initial Memory	Jboss JVM Memory	False
JVM Memory|Heap Memory Usage|Maximum Memory	Jboss JVM Memory	False
JVM Memory|Heap Memory Usage|Used Memory	Jboss JVM Memory	True
JVM Memory|Non Heap Memory Usage|Committed Memory	Jboss JVM Memory	False
JVM Memory|Non Heap Memory Usage|Initial Memory	Jboss JVM Memory	False
JVM Memory|Non Heap Memory Usage|Maximum Memory	Jboss JVM Memory	False
JVM Memory|Non Heap Memory Usage|Used Memory	Jboss JVM Memory	False
JVM Memory|Object Pending Finalization Count	Jboss JVM Memory	True
UTILIZATION|Active Count	Jboss Datasource Pool	False
UTILIZATION|Available Count	Jboss Datasource Pool	False
JVM Memory Pool<InstanceName>|Collection Usage|Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Collection Usage|Initial Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Collection Usage|Used Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Collection Usage|Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Peak Usage|Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Peak Usage|Initial Memory	Jboss JVM Memory Pool	False

Table 1-32. JBoss EAP Metrics (continued)

Metric Name	Category	KPI
JVM Memory Pool<InstanceName>|Peak Usage|Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Peak Usage|Used Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Usage|Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Usage|Initial Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Usage|Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName>|Usage|Used Memory	Jboss JVM Memory Pool	False

RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

Table 1-33. RabbitMQ Metrics

Metric Name	Category	KPI
CPU|Limit	RabbitMQ	False
CPU|Used	RabbitMQ	True
Disk|Free	RabbitMQ	False
Disk|Free limit	RabbitMQ	False
FileDescriptor|Total	RabbitMQ	False
FileDescriptor|Used	RabbitMQ	False
Memory|Limit	RabbitMQ	False
Memory|Used	RabbitMQ	True
Messages|Acked	RabbitMQ	False
Messages|Delivered	RabbitMQ	False
Messages|Delivered get	RabbitMQ	False
Messages|Published	RabbitMQ	False
Messages|Ready	RabbitMQ	False
Messages|Unacked	RabbitMQ	False
Socket|Limit	RabbitMQ	False
Socket|Used	RabbitMQ	True

Table 1-33. RabbitMQ Metrics (continued)

Metric Name	Category	KPI
UTILIZATION|Channels	RabbitMQ	True
UTILIZATION|Connections	RabbitMQ	True
UTILIZATION|Consumers	RabbitMQ	True
UTILIZATION|Exchanges	RabbitMQ	True
UTILIZATION|Messages	RabbitMQ	True
UTILIZATION|Queues	RabbitMQ	True
Messages|Publish in	RabbitMQ Exchange	False
Messages|Publish out	RabbitMQ Exchange	False
Consumer Utilisation	RabbitMQ Queue	False
Consumers	RabbitMQ Queue	False
Memory	RabbitMQ Queue	False
Messages|Ack	RabbitMQ Queue	False
Messages|Ack rate	RabbitMQ Queue	False
Messages|Deliver	RabbitMQ Queue	False
Messages|Deliver get	RabbitMQ Queue	False
Messages|Persist	RabbitMQ Queue	False
Messages|Publish	RabbitMQ Queue	False
Messages|Publish rate	RabbitMQ Queue	False
Messages|Ram	RabbitMQ Queue	False
Messages|Ready	RabbitMQ Queue	False
Messages|Redeliver	RabbitMQ Queue	False
Messages|Redeliver rate	RabbitMQ Queue	False
Messages|Space	RabbitMQ Queue	False
Messages|Unack	RabbitMQ Queue	False
Messages|Unacked	RabbitMQ Queue	False
Messages	RabbitMQ Queue	False

There are no metrics collected for RabbitMQ Virtual Host.

MySQL Metrics

Metrics are collected for the MySQL application service.

Table 1-34. MySQL Metrics

Metric Name	Category	KPI
Aborted connection count	MySQL	True
Connection count	MySQL	True
Event wait average time	MySQL	False
Event wait count	MySQL	False
Binary Files|Binary Files Count	MySQL	False
Binary Files|Binary Size Bytes	MySQL	False
Global Status|Aborted Clients	MySQL	False
Global Status|Binlog Cache Disk Use	MySQL	False
Global Status|Bytes Received	MySQL	False
Global Status|Bytes Sent	MySQL	False
Global Status|Connection Errors Accept	MySQL	False
Global Status|Connection Errors Internal	MySQL	False
Global Status|Connection Errors Max Connections	MySQL	False
Global Status|Queries	MySQL	False
Global Status|Threads Cached	MySQL	False
Global Status|Threads Connected	MySQL	False
Global Status|Threads Running	MySQL	False
Global Status|Uptime	MySQL	False
Global Variables|Delayed Insert Limit	MySQL	False
Global Variables|Delayed Insert Timeout	MySQL	False
Global Variables|Delayed Queue Size	MySQL	False
Global Variables|Max Connect Errors	MySQL	False
Global Variables|Max Connections	MySQL	False
Global Variables|Max Delayed Threads	MySQL	False
Global Variables|Max Error Count	MySQL	False
InnoDB|All deadlock count	MySQL	False

Table 1-34. MySQL Metrics (continued)

Metric Name	Category	KPI
InnoDB|Buffer Pool Bytes Data	MySQL	False
InnoDB|Buffer Pool Bytes Data	MySQL	False
InnoDB|Buffer Pool Bytes Dirty	MySQL	False
InnoDB|Buffer Pool Dump Status	MySQL	False
InnoDB|Buffer Pool Load Status	MySQL	False
InnoDB|Buffer Pool Pages Data	MySQL	False
InnoDB|Buffer Pool Pages Dirty	MySQL	False
InnoDB|Buffer Pool Pages Flushed	MySQL	False
InnoDB|Buffer pool size	MySQL	True
InnoDB|Checksums	MySQL	False
InnoDB|Open file count	MySQL	False
InnoDB|Row lock average time	MySQL	False
InnoDB|Row lock current waits	MySQL	False
InnoDB|Row lock maximum time	MySQL	False
InnoDB|Row lock time	MySQL	False
InnoDB|Row lock waits	MySQL	True
InnoDB|Table lock count	MySQL	False
Performance Table IO Waits|IO Waits Total Delete	MySQL	False
Performance Table IO Waits|IO Waits Total Fetch	MySQL	False
Performance Table IO Waits|IO Waits Total Insert	MySQL	False
Performance Table IO Waits|IO Waits Total Update	MySQL	False
Process List|Connections	MySQL	False
IO waits average time	MySQL Database	False
IO waits count	MySQL Database	True
Read high priority average time	MySQL Database	False
Read high priority count	MySQL Database	False

Table 1-34. MySQL Metrics (continued)

Metric Name	Category	KPI
Write concurrent insert average time	MySQL Database	False
Write concurrent insert count	MySQL Database	False

NGINX Metrics

Metrics are collected for the NGINX application service.

Table 1-35. NGINX Metrics

Metric Name	Category	KPI
HTTP Status Info|Accepts	Nginx	True
HTTP Status Info|Active connections	Nginx	False
HTTP Status Info|Handled	Nginx	True
HTTP Status Info|Reading	Nginx	False
HTTP Status Info|Requests	Nginx	False
HTTP Status Info|Waiting	Nginx	True
HTTP Status Info|Writing	Nginx	False

Sharepoint Metrics

Metrics are collected for the Sharepoint application service.

Table 1-36. Sharepoint Metrics

Metric Name	Category	KPI
Sharepoint Foundation|Active Threads	SharePoint Server	True
Sharepoint Foundation|Current Page Requests	SharePoint Server	False
Sharepoint Foundation|Executing SQL Queries	SharePoint Server	False
Sharepoint Foundation|Executing Time/Page Request	SharePoint Server	True
Sharepoint Foundation|Incoming Page Requests Rate	SharePoint Server	False
Sharepoint Foundation|Object Cache Hit Count	SharePoint Server	False
Sharepoint Foundation|Reject Page Requests Rate	SharePoint Server	False
Sharepoint Foundation|Responded Page Requests Rate	SharePoint Server	True
SQL query executing time	SharePoint Server	False

Table 1-36. Sharepoint Metrics (continued)

Metric Name	Category	KPI
Network|Received Data Rate	SharePoint Web Server	True
Network|Sent Data Rate	SharePoint Web Server	True
Process|Processor Time (%)	SharePoint Windows Service	False
Process|Threads	SharePoint Windows Service	False

Oracle Weblogic Metrics

Metrics are collected for the Oracle Weblogic application service.

Table 1-37. Oracle Weblogic Metrics

Metric Name	Category	KPI
UTILIZATION|Process Cpu Load	Oracle WebLogic Server	True
UTILIZATION|System Cpu Load	Oracle WebLogic Server	False
UTILIZATION|System Load Average	Oracle WebLogic Server	False
UTILIZATION|Collection Time	Weblogic Garbage Collector	True
UTILIZATION|Connections HighCount	Weblogic JMS Runtime	True
UTILIZATION|JMS Servers TotalCount	Weblogic JMS Runtime	False
UTILIZATION|Active Total Count Used	Weblogic JTA Runtime	False
UTILIZATION|Active Transactions TotalCount	Weblogic JTA Runtime	False
UTILIZATION|Transaction Abandoned TotalCount	Weblogic JTA Runtime	True
UTILIZATION|Transaction RolledBack App TotalCount	Weblogic JTA Runtime	True
UTILIZATION|Heap Memory Usage	Weblogic JVM Memory	True
UTILIZATION|Non Heap Memory Usage	Weblogic JVM Memory	False
UTILIZATION|Peak Usage	Weblogic JVM Memory Pool	True
UTILIZATION|Usage	Weblogic JVM Memory Pool	False
UTILIZATION|UpTime	Weblogic JVM Runtime	False

Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

Table 1-38. Pivotal TC Server Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName>|Count	Pivotal TC Server	False
Buffer Pool<InstanceName>|Memory Used	Pivotal TC Server	False
Buffer Pool<InstanceName>|Total Capacity	Pivotal TC Server	False
Class Loading|Loaded Class Count	Pivotal TC Server	False
Class Loading|Total Loaded Class Count	Pivotal TC Server	False
Class Loading|Unloaded Class Count	Pivotal TC Server	False
File Descriptor Usage|Max File Descriptor Count	Pivotal TC Server	False
File Descriptor Usage|Open File Descriptor Count	Pivotal TC Server	False
Garbage Collection:<InstanceName>| Total Collection Count	Pivotal TC Server	False
Garbage Collection:<InstanceName>| Total Collection Time	Pivotal TC Server	False
Process CPU Usage (%)	Pivotal TC Server	True
JVM Memory|Heap Memory Usage| Committed Memory	Pivotal TC Server	True
JVM Memory|Heap Memory Usage| Initial Memory	Pivotal TC Server	False
JVM Memory|Heap Memory Usage| Maximum Memory	Pivotal TC Server	False
JVM Memory|Heap Memory Usage| Used Memory	Pivotal TC Server	True
JVM Memory|Non Heap Memory Usage|Committed Memory	Pivotal TC Server	True
JVM Memory|Non Heap Memory Usage|Initial Memory	Pivotal TC Server	False
JVM Memory|Non Heap Memory Usage|Maximum Memory	Pivotal TC Server	False
JVM Memory|Non Heap Memory Usage|Used Memory	Pivotal TC Server	True
JVM Memory|Number of Object Pending Finalization Count	Pivotal TC Server	True
JVM Memory|Pool:<InstanceName>| Peak Usage|Committed Memory	Pivotal TC Server	False

Table 1-38. Pivotal TC Server Metrics (continued)

Metric Name	Category	KPI
JVM Memory|Pool:<InstanceName>| Peak Usage|Initial Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Peak Usage|Maximum Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Peak Usage|Used Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Usage|Committed Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Usage|Initial Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Usage|Maximum Memory	Pivotal TC Server	False
JVM Memory|Pool:<InstanceName>| Usage|Used Memory	Pivotal TC Server	False
Process CPU Usage (%)	Pivotal TC Server	True
System CPU Usage (%)	Pivotal TC Server	True
Uptime	Pivotal TC Server	True
Threading|Thread Count	Pivotal TC Server	False
System Load Average	Pivotal TC Server	False
Current Thread Count	Pivotal TC Server Thread Pool	False
Current Threads Busy	Pivotal TC Server Thread Pool	True
Total Request Bytes Received	Pivotal TC Server Thread Pool	False
Total Request Bytes Sent	Pivotal TC Server Thread Pool	False
Total Request Count	Pivotal TC Server Thread Pool	True
Total Request Error Count	Pivotal TC Server Thread Pool	True
Total Request Processing Time	Pivotal TC Server Thread Pool	True
JSP Count	Pivotal TC Server Web Module	False
JSP Reload Count	Pivotal TC Server Web Module	False
JSP Unload Count	Pivotal TC Server Web Module	False

ActiveMQ Metrics
Metrics are collected for the ActiveMQ application service.

Table 1-39. ActiveMQ Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName>|Count	Active MQ	False
Buffer Pool<InstanceName>| Memory Used	Active MQ	False
Buffer Pool<InstanceName>|Total Capacity	Active MQ	False
Class Loading|Loaded Class Count	Active MQ	False
Class Loading|Unloaded Class Count	Active MQ	False
Class Loading|Total Loaded Class Count	Active MQ	False
File Descriptor Usage|Max File Descriptor Count	Active MQ	False
File Descriptor Usage|Open File Descriptor Count	Active MQ	False
Garbage Collection<InstanceName>| Total Collection Count	Active MQ	False
Garbage Collection<InstanceName>| Total Collection Time	Active MQ	False
JVM Memory Pool<InstanceName>|Peak Usage|Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName>|Peak Usage|Initial Memory	Active MQ	False
JVM Memory Pool<InstanceName>|Peak Usage|Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName>|Peak Usage|Used Memory	Active MQ	False
JVM Memory Pool<InstanceName>| Usage|Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName>| Usage|Initial Memory	Active MQ	False

Table 1-39. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
JVM Memory Pool<InstanceName>|Usage|Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName>|Usage|Used Memory	Active MQ	False
Threading|Thread Count	Active MQ	False
Uptime	Active MQ	False
UTILIZATION|Process CpuLoad	Active MQ	False
UTILIZATION|Memory Limit	ActiveMQ Broker	True
UTILIZATION|Memory Percent Usage (%)	ActiveMQ Broker	True
UTILIZATION|Store Limit	ActiveMQ Broker	False
UTILIZATION|Store Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION|Temp Limit	ActiveMQ Broker	False
UTILIZATION|Temp Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION|Total Consumer Count	ActiveMQ Broker	True
UTILIZATION|Total Dequeue Count	ActiveMQ Broker	True
UTILIZATION|Total Enqueue Count	ActiveMQ Broker	True
UTILIZATION|Total Message Count	ActiveMQ Broker	True
JVM Memory|Heap Memory Usage|Initial Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Heap Memory Usage|Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Heap Memory Usage|Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Heap Memory Usage|Used Memory	ActiveMQ JVM Memory Usage	False

Table 1-39. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
JVM Memory|Non Heap Memory Usage|Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Non Heap Memory Usage|Initial Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Non Heap Memory Usage|Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Non Heap Memory Usage|Used Memory	ActiveMQ JVM Memory Usage	False
JVM Memory|Object Pending FinalizationCount	ActiveMQ JVM Memory Usage	False
UTILIZATION|Process CpuLoad	ActiveMQ OS	False
UTILIZATION|System Cpu Load	ActiveMQ OS	False
UTILIZATION|Consumer Count	ActiveMQ Topic	True
UTILIZATION|Dequeue Count	ActiveMQ Topic	True
UTILIZATION|Enqueue Count	ActiveMQ Topic	True
UTILIZATION|Queue Size	ActiveMQ Topic	True
UTILIZATION|Producer Count	ActiveMQ Topic	False

Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

Note Metrics are collected for the Events MPM. Metrics are not collected for the other MPMs.

Table 1-40. Apache HTTPD Metrics

Metric Name	Category	KPI
UTILIZATION|Busy Workers	Apache HTTPD	True
UTILIZATION|Bytes Per Req	Apache HTTPD	False
UTILIZATION|Bytes Per Sec	Apache HTTPD	False
UTILIZATION|CPU Load	Apache HTTPD	True
UTILIZATION|CPU User	Apache HTTPD	False

Table 1-40. Apache HTTPD Metrics (continued)

Metric Name	Category	KPI
UTILIZATION|Idle Workers	Apache HTTPD	True
UTILIZATION|Request Per Sec	Apache HTTPD	True
UTILIZATION|SCBoard Closing	Apache HTTPD	False
UTILIZATION|SCBoard DNS Lookup	Apache HTTPD	False
UTILIZATION|SCBoard Finishing	Apache HTTPD	False
UTILIZATION|SCBoard Idle Cleanup	Apache HTTPD	False
UTILIZATION|SCBoard Keep Alive	Apache HTTPD	False
UTILIZATION|SCBoard Logging	Apache HTTPD	False
UTILIZATION|SCBoard Open	Apache HTTPD	False
UTILIZATION|SCBoard Reading	Apache HTTPD	False
UTILIZATION|SCBoard Sending	Apache HTTPD	False
UTILIZATION|SCBoard Starting	Apache HTTPD	False
UTILIZATION|SCBoard Waiting	Apache HTTPD	False
UTILIZATION|Total Accesses	Apache HTTPD	False
UTILIZATION|Total Bytes	Apache HTTPD	True
UTILIZATION|Total Connections	Apache HTTPD	False
UTILIZATION|Uptime	Apache HTTPD	True
UTILIZATION|Asynchronous Closing Connections	Apache HTTPD	False
UTILIZATION|Asynchronous Keep Alive Connections	Apache HTTPD	False
UTILIZATION|Asynchronous Writing Connections	Apache HTTPD	False
UTILIZATION|ServerUptimeSeconds	Apache HTTPD	False
UTILIZATION|Load1	Apache HTTPD	False
UTILIZATION|Load5	Apache HTTPD	False
UTILIZATION| ParentServerConfigGeneration	Apache HTTPD	False
UTILIZATION| ParentServerMPMGeneration	Apache HTTPD	False

Oracle Database Metrics

Metrics are collected for the Oracle database application service.

Oracle database cannot be activated on Linux platforms.

Table 1-41. Oracle Database Metrics

Metric Name	Category	KPI
Utilization|Active Sessions	OracleDB	True
Utilization|Buffer CacheHit Ratio	OracleDB	False
Utilization|Cursor CacheHit Ratio	OracleDB	False
Utilization|Database Wait Time	OracleDB	False
Utilization|Disk Sort persec	OracleDB	False
Utilization|Enqueue Timeouts Persec	OracleDB	False
Utilization|Global Cache Blocks Corrupted	OracleDB	False
Utilization|Global Cache Blocks Lost	OracleDB	False
Utilization|Library CacheHit Ratio	OracleDB	False
Utilization|Logon persec	OracleDB	True
Utilization|Memory Sorts Ratio	OracleDB	True
Utilization|Rows persort	OracleDB	False
Utilization|Service Response Time	OracleDB	False
Utilization|Session Count	OracleDB	True
Utilization|Session Limit	OracleDB	False
Utilization|Shared Pool Free	OracleDB	False
Utilization|Temp Space Used	OracleDB	False
Utilization|Total Sorts persec	OracleDB	False
Utilization|Physical Read Bytes Persc	OracleDB	False
Utilization|Physical Read IO Requests Persc	OracleDB	False
Utilization|Physical Read Total Bytes Persec	OracleDB	False
Utilization|Physical Reads Persec	OracleDB	True
Utilization|Physical Reads Per Txn	OracleDB	False
Utilization|Physical Write Bytes Persc	OracleDB	False
Utilization|Physical Write IO Requests Persc	OracleDB	False

Table 1-41. Oracle Database Metrics (continued)

Metric Name	Category	KPI
Utilization|Physical Write Total Bytes Persc	OracleDB	False
Utilization|Physical Writes Persc	OracleDB	True
Utilization|Physical Writes Per Txn	OracleDB	False
Utilization|User Commits Percentage	OracleDB	False
Utilization|User Commits Persc	OracleDB	False
Utilization|User Rollbacks Percentage	OracleDB	False
Utilization|User Rollbacks persec	OracleDB	True
Utilization|User Transaction Persec	OracleDB	False
Utilization|Database Time Persc	OracleDB	False

Cassandra Database Metrics

Metrics are collected for the Cassandra database application service.

Table 1-42. Cassandra Database Metrics

Metric Name	Category	KPI
Cache<InstanceName>|Capacity	Cassandra	False
Cache<InstanceName>|Entries	Cassandra	True
Cache<InstanceName>|HitRate	Cassandra	True
Cache<InstanceName>|Requests	Cassandra	True
Cache<InstanceName>|Size	Cassandra	False
ClientRequest<InstanceName>|Failures	Cassandra	False
ClientRequest<InstanceName>|Latency	Cassandra	False
ClientRequest<InstanceName>|Timeouts	Cassandra	False
ClientRequest<InstanceName>|Total Latency	Cassandra	False
ClientRequest<InstanceName>|Unavailables	Cassandra	False
CommitLog|Pending Tasks	Cassandra	False
CommitLog|Total Commit Log Size	Cassandra	False
Compaction|Bytes Compacted	Cassandra	False

Table 1-42. Cassandra Database Metrics (continued)

Metric Name	Category	KPI
Compaction|Completed Tasks	Cassandra	False
Compaction|Pending Tasks	Cassandra	False
Compaction|Total Compactions Completed	Cassandra	False
Connected Native Clients	Cassandra	False
HeapMemoryUsage|committed	Cassandra	False
HeapMemoryUsage|init	Cassandra	False
HeapMemoryUsage|max	Cassandra	False
HeapMemoryUsage|used	Cassandra	False
NonHeapMemoryUsage|committed	Cassandra	False
NonHeapMemoryUsage|init	Cassandra	False
NonHeapMemoryUsage|max	Cassandra	False
NonHeapMemoryUsage|used	Cassandra	False
ObjectPendingFinalizationCount	Cassandra	False
Storage|Exceptions Count	Cassandra	False
Storage|Load Count	Cassandra	False
Table<InstanceName>|Coordinator Read Latency	Cassandra	False
Table<InstanceName>|Live Diskspace Used	Cassandra	False
Table<InstanceName>|Read Latency	Cassandra	False
Table<InstanceName>|Total Diskspace Used	Cassandra	False
Table<InstanceName>|Total Read Latency	Cassandra	False
Table<InstanceName>|Total Write Latency	Cassandra	False
Table<InstanceName>|Write Latency	Cassandra	False
ThreadPools<InstanceName>|Active Tasks	Cassandra	False

Table 1-42. Cassandra Database Metrics (continued)

Metric Name	Category	KPI
ThreadPools<InstanceName>| Currently Blocked Tasks	Cassandra	False
ThreadPools<InstanceName>| Pending Tasks	Cassandra	False

Hyper-V Metrics
Metrics are collected for the Hyper-V application service.

Table 1-43. Hyper-V Metrics

Metric Name	Category	KPI
VM:Hyper-V Virtual Machine Health Summary|Health Critical	HyperV	False
VM<instanceName>|Physical Memory	HyperV	False
VM<instanceName>Hv VP 0|Total Run Time	HyperV	False
VM<instanceName>|Bytes Received	HyperV	False
VM<instanceName>|Bytes Sent	HyperV	False
VM<instanceName>|Error Count	HyperV	False
VM<instanceName>|Latency	HyperV	False
VM<instanceName>|Queue Length	HyperV	False
VM<instanceName>|Throughput	HyperV	False
CPU<instanceName>|Idle Time	HyperV	True
CPU<instanceName>|Processor Time	HyperV	True
CPU<instanceName>|User Time	HyperV	True
Disk<instanceName>|Avg Disk Queue Length	HyperV	False
Disk<instanceName>|Idle Time	HyperV	False
Disk<instanceName>|Read Time	HyperV	True
Disk<instanceName>|Write Time	HyperV	True
Process<instanceName>|Private Bytes	HyperV	False
Process<instanceName>|Processor Time	HyperV	False
Process<instanceName>|Thread Count	HyperV	False
Process<instanceName>|User Time	HyperV	False

Table 1-43. Hyper-V Metrics (continued)

Metric Name	Category	KPI
System|Processes	HyperV	False
System|Processor Queue Length	HyperV	False
System|System UpTime	HyperV	False
Memory|Available Bytes	HyperV	False
Memory|Cache Bytes	HyperV	False
Memory|Cache Faults	HyperV	False
Memory|Pages	HyperV	False
Network<instanceName>|Packets Outbound Error	HyperV	False
Network<instanceName>|Packets Received Error	HyperV	False

MongoDB Metrics

Metrics are collected for the MongoDB application service.

Table 1-44. MongoDB Metrics

Metric Name	Category	KPI
UTILIZATION|Active Reads	MongoDB	True
UTILIZATION|Active Writes	MongoDB	True
UTILIZATION|Connections Available	MongoDB	False
UTILIZATION|Connections Total Created	MongoDB	False
UTILIZATION|Current Connections	MongoDB	True
UTILIZATION|Cursor Timed Out	MongoDB	True
UTILIZATION|Deletes Per Sec	MongoDB	False
UTILIZATION|Document Inserted	MongoDB	False
UTILIZATION|Document Deleted	MongoDB	False
UTILIZATION|Flushes Per Sec	MongoDB	False
UTILIZATION|Inserts Per Sec	MongoDB	False
UTILIZATION|Net Input Bytes	MongoDB	False
UTILIZATION|Open Connections	MongoDB	True
UTILIZATION|Page Faults Per Second	MongoDB	False

Table 1-44. MongoDB Metrics (continued)

Metric Name	Category	KPI
UTILIZATION|Net Output Bytes	MongoDB	False
UTILIZATION|Queries Per Sec	MongoDB	False
UTILIZATION|Queued Reads	MongoDB	True
UTILIZATION|Queued Writes	MongoDB	True
UTILIZATION|Total Available	MongoDB	False
UTILIZATION|Total Deletes Per Sec	MongoDB	False
UTILIZATION|Total Passes Per Sec	MongoDB	False
UTILIZATION|Total Refreshing	MongoDB	False
UTILIZATION|Updates Per Sec	MongoDB	False
UTILIZATION|Volume Size MB	MongoDB	False
UTILIZATION|Collection Stats	MongoDB DataBases	False
UTILIZATION|Data Index Stats	MongoDB DataBases	True
UTILIZATION|Data Indexes	MongoDB DataBases	False
UTILIZATION|Data Size Stats	MongoDB DataBases	True
UTILIZATION|Average Object Size stats	MongoDB DataBases	False
UTILIZATION|Num Extents Stats	MongoDB DataBases	False

Riak Metrics

Metrics are collected for the Riak application service.

Table 1-45. Riak Metrics

Metric Name	Category	KPI
UTILIZATION|CPU Average	Riak KV	False
UTILIZATION|Memory Processes	Riak KV	False
UTILIZATION|Memory Total	Riak KV	False
UTILIZATION|Node GETs	Riak KV	True
UTILIZATION|Node GETs Total	Riak KV	False
UTILIZATION|Node PUTs	Riak KV	True
UTILIZATION|Node PUTs Total	Riak KV	False
UTILIZATION|PBC Active	Riak KV	True

Table 1-45. Riak Metrics (continued)

Metric Name	Category	KPI
UTILIZATION|PBC Connects	Riak KV	True
UTILIZATION|Read Repairs	Riak KV	True
UTILIZATION|vNODE Index Reads	Riak KV	True
UTILIZATION|vNODE Index Writes	Riak KV	True

NTPD Metrics

Metrics are collected for the NTPD application service.

Table 1-46. NTPD Metrics

Metric Name	Category	KPI
ntpd |delay	Network Time Protocol	True
ntpd | jitter	Network Time Protocol	True
ntpd | offset	Network Time Protocol	True
ntpd | poll	Network Time Protocol	False
ntpd | reach	Network Time Protocol	True
ntpd | when	Network Time Protocol	False

WebSphere Metrics

Metrics are collected for the WebSphere application service.

Table 1-47. WebSphere Metrics

Metric Name	Category	KPI
Thread Pool|Active Count|Current	Thread Pool	False
Thread Pool|Active Count|High	Thread Pool	False
Thread Pool|Active Count|Low	Thread Pool	False
Thread Pool|Active Count|Lower	Thread Pool	False
Thread Pool|Active Count|Upper	Thread Pool	False
JDBC|Close Count	JDBC	False
JDBC|Create Count	JDBC	False
JDBC|JDBC Pool Size|Average	JDBC	False

Table 1-47. WebSphere Metrics (continued)

Metric Name	Category	KPI
JDBC|JDBC Pool Size| Current	JDBC	False
JDBC|JDBC Pool Size| Lower	JDBC	False
JDBC|JDBC Pool Size| Upper	JDBC	False
Garbage Collection<InstanceName>| Total Collection Count	WebSphere	False
Garbage Collection<InstanceName>| Total Collection Time	WebSphere	False
JVM Memory|Heap Memory Usage|Committed Memory	WebSphere	False
JVM Memory|Heap Memory Usage|Initial Memory	WebSphere	False
JVM Memory|Heap Memory Usage|Maximum Memory	WebSphere	False
JVM Memory|Heap Memory Usage|Used Memory	WebSphere	False
JVM Memory|Non Heap Memory Usage|Committed Memory	WebSphere	False
JVM Memory|Non Heap Memory Usage|Initial Memory	WebSphere	False
JVM Memory|Non Heap Memory Usage|Maximum Memory	WebSphere	False
JVM Memory|Non Heap Memory Usage|Used Memory	WebSphere	False
JVM Memory|Number of Object Pending Finalization Count	WebSphere	False
JVM Memory| Pool<InstanceName>|Peak Usage|Committed Memory	WebSphere	False

Table 1-47. WebSphere Metrics (continued)

Metric Name	Category	KPI
JVM Memory|Pool<InstanceName>|Peak Usage|Initial Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Peak Usage|Maximum Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Peak Usage|Used Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Usage|Committed Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Usage|Initial Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Usage|Maximum Memory	WebSphere	False
JVM Memory|Pool<InstanceName>|Usage|Used Memory	WebSphere	False
Process Cpu Load	WebSphere	False
System Cpu Load	WebSphere	False
System Load Average	WebSphere	False

Java Application Metrics

Metrics are collected for the Java application service.

Table 1-48. Java Application Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName>|Count	Java Application	False
Buffer Pool<InstanceName>|Memory Used	Java Application	False
Buffer Pool<InstanceName>|Total Capacity	Java Application	False
Class Loading|Loaded Class Count	Java Application	True
Class Loading|Total Loaded Class Count	Java Application	False
Class Loading|Unloaded Class Count	Java Application	False
Garbage Collection<InstanceName>|Total Collection Count	Java Application	False

Table 1-48. Java Application Metrics (continued)

Metric Name	Category	KPI
Garbage Collection<InstanceName>|Total Collection Time	Java Application	False
JVM Memory|Heap Memory Usage|Committed Memory	Java Application	False
JVM Memory|Heap Memory Usage|Initial Memory	Java Application	False
JVM Memory|Heap Memory Usage|Maximum Memory	Java Application	False
JVM Memory|Heap Memory Usage|Used Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Peak Usage|Committed Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Peak Usage|Initial Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Peak Usage|Maximum Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Peak Usage|Used Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Usage|Committed Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Usage|Initial Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Usage|Maximum Memory	Java Application	False
JVM Memory|JVM Memory Pool<InstanceName>|Usage|Used Memory	Java Application	False
JVM Memory|Non Heap Memory Usage|Committed Memory	Java Application	False
JVM Memory|Non Heap Memory Usage|Initial Memory	Java Application	False
JVM Memory|Non Heap Memory Usage|Maximum Memory	Java Application	False
JVM Memory|Non Heap Memory Usage|Used Memory	Java Application	False

Table 1-48. Java Application Metrics (continued)

Metric Name	Category	KPI
JVM Memory|Object Pending Finalization Count	Java Application	False
Uptime	Java Application	True
Threading|Thread Count	Java Application	True
Process CPU Usage %	Java Application	False
System CPU Usage %	Java Application	False
System Load Average %	Java Application	False

Remote Check Metrics

Metrics are collected for object types such as HTTP, ICMP, TCP, and UDP.

HTTP Metrics

vRealize Operations Manager discovers metrics for HTTP remote checks.

HTTP Metrics

Table 1-49. HTTP Metrics

Metric Name	KPI
Availability	False
Response Code	False
Response Time	True
Result Code	False

ICMP Metrics

vRealize Operations Manager discovers metrics for the ICMP object type.

Table 1-50. ICMP Metrics

Metric Name	KPI
Availability	False
Average Response Time	True
Packet Loss (%)	False
Packets Received	False
Packets Transmitted	False
Result Code	False

TCP Metrics

vRealize Operations Manager discovers metrics for the TCP object type.

Table 1-51. TCP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

UDP Metrics

vRealize Operations Manager discovers metrics for the UDP object type.

Table 1-52. UDP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

Troubleshooting

Troubleshooting the Configuration of vRealize Application Remote Collector

vRealize Application Remote Collector Configuration Fails

An error occurs when you add a vCenter Server while configuring the vRealize Application Remote Collector.

Problem

Configuration of vRealize Application Remote Collector fails with the following error:

```
Unable to establish a valid connection to the target system.
Wait for response of Task 'Test connection' is timed out for collector
'vRealize Operations Manager Collector-Master'.
```

Solution

◆ Enable the relevant ports.

◆ Ensure that vRealize Operations Manager and vRealize Application Remote Collector have the NTP synced.

Error in Password

An error occurs when a colon character is added to the password while deploying vRealize Application Remote Collector.

Problem

Configuration of vRealize Application Remote Collector fails with the following error:

```
Unable to establish a valid connection to the target system. An internal server error has
occurred. Please ensure that the Application Remote Collector is setup properly. For more
details, refer the troubleshooting section in Application Remote Collector product
document.
```

Solution

◆ Ensure that you do not use a colon character : in the `API Admin User's Password` text box while deploying vRealize Application Remote Collector.

Troubleshooting Agent Installation

Agent Install Failure Because of the vCenter Server User Permissions

Guest operation privileges are required to install agents on virtual machines.

Problem

Agent installation fails with the following error message if there are no guest operation privileges:

```
vCenter adapter user is missing either of the following guest operations privileges –
execute, modify, query
```

Solution

1 Verify that you have configured a vCenter adapter.

2 The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: **Guest operation modifications**, **Guest operation program execution**, and **Guest operation queries**.

Agent Install Failure Because NTP is Not in Sync

If the actual time of the vRealize Application Remote Collector server is behind or ahead of the current time, you might face configuration or installation failures.

Problem

- Agent installation fails

- Adapter configuration fails

Solution

◆ Ensure that you configure network time protocol settings, or

◆ Run the following command to update the time immediately from an NTP server: `ntpdate time.vmware.com`

  Ensure that you have stopped the ntpd service before you run the `ntpdate` command.

Note The system time takes about five minutes to sync with the NTP server time.

Agent Install Fails on a Linux End Point

Install of an agent on a Linux end point fails for a non-root user with a specific set of privileges.

Problem

Agent installation fails with the following error if the `tty` command is not added:

```
Bootstrap Failed for VM <VM ID> with error message:{ "status":"FAILED", "data":
[ { "status":"FAILED", "message":"Failed - install - passwordless sudo access is required for
the user <Install Username> on the command mkdir. [sudo: sorry, you must have a tty to run
sudo]", "stage":"0" } ], "currentstage":"0", "totalstages":"0" }
```

Solution

◆ If you get the error as stated above, verify that the following lines exist in `/etc/sudoers`.

```
1. root ALL=(ALL:ALL) ALL
2.Defaults:root !requiretty
3.Defaults:arcuser !requiretty
```

  (1) can be omitted if password-less sudo is already enabled for the root user. (2) and (3) can be omitted if your endpoint VMs are already configured to turn off `requiretty`.

  Add these lines to `/etc/sudoers`, if you have not added them.

◆ To solve other failures on Linux end points, ensure that `/tmp` mount point is mounted with the exec mount option.

Agent Install on Windows Fails When UAC is Disabled

Problem

Install of the agent fails even when UAC is disabled.

Solution

◆ To disable UAC (previously known as LUA) on Windows, complete the following steps:

  a In the registry path
  `HKLM:\Software\Microsoft\Windows\CurrentVersion\Policies\System`, set the value for the key `EnableLUA` to **0**.

  b You must reboot the machine for the changes to take effect.

Agent Install Fails on Windows with a Permission Denied Error

In Windows, during bootstrap, when the Telegraf folder is renamed to `ucp-telegraf`, it can result in a failure because of a permission error.

Problem

Sometimes, there are certain antiviruses running, which prevent the application from renaming or modifying the directory or files. In such a situation, the following error message is displayed:

```
Install telegraf [unable to install telegraf due to system error : [WinError
5] Access is denied: 'C:\\VMware\\UCP\\ucp-telegraf']"}].
```

Solution

◆ Disable the antivirus and then proceed with bootstrapping.

Troubleshooting Plugin Related Failures

Unable to Activate a Plugin

Unable to activate a plugin with the same fields until the plugin configuration is deleted.

Problem

An error message is displayed in the user interface of vRealize Operations Manager that states the following:

```
Failed to update resource: Resource with same key already exists
```

Solution

◆ Manually delete the existing plugin configuration and then continue with the activation of the plugin. If the problem persists, delete the corresponding resource from the inventory.

Plugin Status Is Displayed as Unknown

The status of a few plugins is Unknown after vRealize Application Remote Collector and vRealize Operations Manager are upgraded from 7.5 to 8.2 and 8.1 to 8.2.

Problem

The **Unknown** icon is displayed with a gray icon against the plugin.

Solution

◆ Reactivate the plugin.

Troubleshooting Metric Collection

Troubleshoot Agent Installation and Metric Collection Issues

If the time settings between vRealize Application Remote Collector and vRealize Operations Manager are not synchronized, you might face agent installation and metric collection issues. Eventually, you might not see any metrics in the vRealize Operations Manager dashboards.

Problem

You might notice the following issues in vRealize Operations Manager:

- You cannot add vRealize Application Remote Collector to vRealize Operations Manager

- You cannot install an agent in the Windows and Linux target VMs.

Cause

Time synchronization is a prerequisite of the TLS/SSO communication between client and server.

If the vRealize Operations Manager and vRealize Application Remote Collector are not time synchronized, the test connection fails while configuring vRealize Application Remote Collector in vRealize Operations Manager.

If the Windows and Linux target VMs are not time synchronized with vRealize Operations Manager, communication between vRealize Application Remote Collector and agents will break after installing the agents. Hence monitored metrics are not sent to vRealize Operations Manager. Alternatively, stop and restart the agent to resolve this issue.

Solution

1   Check the vRealize Operations Manager support bundle in the following path: `COLLECTOR/adapters/APPOSUCPAdapter/` for errors.

2   Check the vRealize Application Remote Collector support bundle, *ucpapi.log*, for errors.

3   Ensure time synchronization between vRealize Application Remote Collector, vRealize Operations Manager and the Windows and Linux target VMs.

4   To start and restart the agent, see Additional Operations from the Manage Agents Tab.

Troubleshooting Upgrade

You might see error messages, or inconsistent status icons in vRealize Operations Manager if you do not upgrade to the compatible versions of vRealize Operations Manager and vRealize Application Remote Collector.

Problem

vRealize Application Remote Collector UI Problems

- You cannot update your endpoint VM to have the latest vRealize Application Remote Collector agent.

- If you bootstrap/re-bootstrap a VM after upgrading vRealize Application Remote Collector you cannot activate the newly discovered application. You see an error message if you try to activate it.

Manage vRealize Application Remote Collector UI Problems

- You can see an option to update the endpoint agent but you are unable to perform the update.

- Services supported in the latest versions of vRealize Application Remote Collector cannot be discovered.

Cause

The first set of problems occurs because vRealize Application Remote Collector is upgraded to the latest version, but vRealize Operations Manager is an old version.

The second set of problems occurs because vRealize Operations Manager is upgraded to the latest version, but vRealize Application Remote Collector is in version 1.x.

Solution

- Upgrade to the compatible versions of vRealize Operations Manager and vRealize Application Remote Collector.

    Note For more troubleshooting steps, see Troubleshooting Agent Installation.

Troubleshooting Content Upgrade

Problem

Content upgrade for an end point fails with the following error:

```
Timeout Error. Please retry the action after some time.
```

Cause

Sometimes content upgrade for an end point fails because of a timeout in the vRealize Application Remote Collector.

Solution

- Retrigger content upgrade for the end point to resolve the issue.

Troubleshooting Using Support Bundles

Download the support bundles from the virtual machines where you deployed vRealize Application Remote Collector. Support bundles are required to troubleshoot problems related to application monitoring. For Linux and Windows end point VMs, run the specified command and access the support bundle.

For vRealize Application Remote Collector

1   Access the VAMI page by entering https://<vRealize Application Remote Collector hostname>:5480

2   Log in with root credentials.

3   Click the **Support Bundle** tab. Click the **Generate Logs for VA** button.

     vRealize Application Remote Collector creates the support bundles which you can download.

For End Point VMs

1   Log in to the end point.

2   Run the following commands based on the end point VM's operating system type:

     For Linux End Point VMs

```
/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh --config /opt/vmware/ucp/salt-minion/etc/salt/
grains --action gen_support_bundle --log_level INFO
```

The support bundle is generated and placed as a ZIP file in the `/opt/vmware/ucp/support-bundle-endpoints/` directory.

     For Windows End Point VMs

```
C:\VMware\UCP\ucp-minion\bin\ucp-minion.bat --config C:\VMware\UCP\salt\conf\grains --
action gen_support_bundle --log_level INFO
```

The support bundle is generated and placed as a ZIP file in the `%SystemDrive%\VMware\UCP\support-bundle-endpoints\` directory.

Service Discovery

Service discovery helps you discover services running in each VM and then builds a relationship or dependency between the services from different VMs. You can view basic metrics based on the services you want to monitor. You can also use the service discovery dashboards to monitor the services.

Service discovery helps you determine the kind of services running on each VM in your environment. You can find out which VM is a part of a service, the impact of shutting down or moving a VM, the impact of an incident, and the right escalation path for a problem. You can also determine which VMs are used to migrate a service and which services are impacted by a planned outage on a VM or an infrastructure component.

Licensing

You can discover and monitor services using vRealize Operations Manager Advanced and Enterprise editions.

To discover and monitor services, follow these steps in vRealize Operations Manager:

- Configure Service Discovery. For more information, see Configure Service Discovery.

- Manage Services. For more information, see Manage Services .

- Monitor services using dashboards. For more information, see Service Discovery Dashboards.

- View the services discovered. For more information, see Discovered Services.

Supported Platforms and Products for Service Discovery

Service discovery supports specific platforms and product versions.

You can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

Supported Product Versions for Credential-Based Service Discovery

- For ESXi, vCenter Server, and VMware Cloud on AWS versions, see the VMware Product Interoperability Matrix.

- VMware Tools: For details, see KB 75122.

Supported Product Versions and Other Pre-Requisites for Credential-Less Service Discovery

For information, see KB 78216.

Operating System Versions

Operating Systems	Version
Windows	Windows 7, Windows Server 2008/R2, and above.
Linux	Photon, RHEL, CentOS, SUSE Linux Enterprise Server, OEL, and Ubuntu (all Linux operating systems must be based on kernel version 2.6.25 or above).

Supported Services

Service discovery supports several services that are supported in vRealize Operations Manager. The supported services are listed here.

Supported Services:

- Active Directory

- Apache HTTP

- Apache Tomcat

- DB2

- Exchange Client Access Server

- Exchange Edge Transport Server

- Exchange Hub Transport Server

- Exchange Mailbox Server

- Exchange Server

- Exchange Unified Messaging Server

- GemFire

- IIS

- JBoss

- MS SQL DB

- MySQL DB

- Nginx

- Oracle DB

- RabbitMQ

- SharePoint

- SharePoint Application Server

- SharePoint Server

- SharePoint Web Server

- SRM vCenter Replication Management Server

- SRM vCenter Replication Server

- Sybase DB

- Pivotal tc Server

- vCenter Site Recovery Manager Server

- vCloud Director

- VMware vCenter

- VMware vCenter (Appliance)

- VMware View Server

- vRealize Operations Analytics

- vRealize Operations Collector

- vRealize Operations GemFire

- vRealize Operations Postgres Data

- vRealize Operations Postgres Repl

- vRealize Operations UI

- WebLogic

- WebSphere

Configure Service Discovery

To discover services and their relationships and to access basic monitoring, you can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

Prerequisites

- You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to discover services. The configured vCenter Server user must have the following privileges:

  - Guest operation alias modification

  - Guest operation alias query

  - Guest operation modifications

  - Guest operation program execution

  - Guest operation queries

  - Manage service configurations

  - Modify service configuration

  - Query service configuration

  - Read service configuration

- The ESXi instance that hosts the VMs where services should be discovered, must have HTTPS access to port 443 from the collector node on which the service discovery adapter instance is configured.

- Verify that the following types of commands and utilities are used:

Type	Commands and Utilities
UNIX Operating Systems	
Service Discovery	`ps`, `netstat`, and `top`
Performance Metrics Collection	: `awk`, `csh`, `ps`, `pgrep`, and `procfs` (file system)
Windows Operating Systems	
Service Discovery	`wmic` and `netstat`
Performance Metrics Collection	`wimic`, `typeperf`, and `tasklist`

- User Access Restrictions

  - For Linux operating systems, ensure that the user is a root or member of the *sudo* users group.

    Note For non-root users, the `NOPASSWD` option must be enabled in `/etc/sudoers` file to avoid the metrics collector scripts from waiting for the interactive password input.

    Steps to enable the `NOPASSWD` option for a particular sudo user:

    1 Login to the specific VM as a root user.

    2 Run the `sudo visudo` command that opens an editor.

    3 In the command section, add `username ALL=(ALL) NOPASSWD:ALL`. `username` must be replaced with an existing user name for which this option is enabled.

    4 Save the file and close it. It is automatically reloaded.

  - To discover services on Windows, the local administrator account must be configured.

    Note Services will not be discovered for administrator group members that are different from the administrator account itself if the policy setting `User Account Control: Run all administrators in Admin Approval Mode` is turned on. As a workaround, you can turn off this policy setting to discover services. However, if you turn the policy setting off, the security of the operating system is reduced.

  - To discover services on Windows Active Directory, the domain administrator account must be configured.

- The system clock must be synchronized between the vRealize Operations Manager nodes, the vCenter Server, and the VM if service discovery is working in credential-based mode and guest alias mapping is used for authentication.

- The configured user must have read and write privileges to the temp directory. For Windows systems, the path can be taken from the environment variable *TEMP*. For Linux systems, it is `/tmp` and/or `/var/tmp`.

- For more information about supported platforms and versions, see Supported Platforms and Products for Service Discovery.

Note If more than one vRealize Operations Manager instance is monitoring the same vCenter Server and service discovery is enabled for those vRealize Operations Manager instances, then service discovery might be unstable, which is a known VMware Tools problem. As a result, guest operations might fail to execute.

Procedure

1 In the menu, select **Home** and then select **Manage Applications > Discover Services** from the left panel.

2 From the **Discover Services** page, click the **Configure Service Discovery** option.

3 From the **Cloud Accounts** page, click the vCenter Server instance from the list and then select the **Service Discovery** tab.

4 To enable service discovery in this vCenter Server, enable the **Service Discovery** option.

5 You can choose to add credentials by selecting the **Use alternate credentials** check box.

  a Click the plus sign and enter the details in the **Manage Credentials** dialog box, which include a credential name and a vCenter user name and password. In addition, enter the user name and password for Windows, Linux, and SRM and click **OK**.

6 Alternatively, if you are using the default user name and password, enter a default user name and password for Windows, Linux, and SRM.

7 Enter a password for the guest user mapping.

8 You can also enable grouping of the application and the creation of a business application.

9 Click **Save**.

10 Edit the cloud account created for service discovery.

11 In the **Advanced Settings** section, to configure credential-less service discovery, select **Enabled** from the **Credential-less service discovery status** field.

What to do next

You can manage services supported by vRealize Operations Manager on specific VMs.

Manage Services

You can manage services supported by vRealize Operations Manager on the specific VMs.

Where You Manage Services

In the menu, select **Administration** and then select **Inventory** from the left panel. Select the **Manage Services** tab from the right pane. You can also navigate to the `Manage Services` tab by selecting **Home**, and then select **Manage Applications > Discover Services** from the left pane. Select the `Manage Services` option from the **Discover Services** page.

You can view specific details from the options in the data grid.

Table 1-53. Datagrid Options

Options	Description
VM Name	Name of the VM.
Operating System	Operating system installed on the VM.
Services Discovered	Displays the names of discovered services or `None`, if services are not discovered on the VM.
Service Monitoring	Displays the current value of the VM's service monitoring setting. If set, services are discovered and service performance metrics are calculated every 5 minutes. Otherwise, service discovery is performed every 24 hours.

Table 1-53. Datagrid Options (continued)

Options	Description
Authentications Status	VM authentication status for service discovery. The possible values are: ■ Unknown ■ Failed ■ Guest Alias ■ Common Credentials ■ Credential-less
Power State	Power status of the VMs. The possible values are: ■ Powered On ■ Powered Off ■ Suspended ■ Unknown
Collection State	Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the collection state icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**.
Collection Status	Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the collection status icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**. You can view a message for VMs with a failed authentication status in a tool tip when you point to the collection status icon.
vCenter Name	Name of the vCenter Adapter instance to which that VM resource belongs.

Table 1-54. Toolbar Options

Options	Description
Actions	Displays a list of actions. For more information, see List of vRealize Operations Manager Actions.
Open in external application	If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application.
Provide Password	Select VMs from the list and click **Provide Password** to provide a user name and password for the selected VMs to discover the services.

Table 1-54. Toolbar Options (continued)

Options	Description
Enable Service Monitoring	Select VMs from the list and click **Enable Service Monitoring** to enable frequent service discovery and service performance metrics calculation (every 5 minutes).
	Note Selecting too many VMs will potentially result in vCenter Server degradation which is a known issue.
Disable Service Monitoring	Select VMs from the list and click **Disable Service Monitoring** to disable frequent service discovery and service performance metrics calculation. Service discovery defaults to the 24-hour cycle.
Clear Selections	Clears all VM object selections.
Select All	Selects all VM objects.
Show Detail	Navigates to the **Summary** tab for the selected VM.
Page Size	The number of objects to list per page.
All Filters	You can search through the list of VMs according to the following criteria: VM Name, Operating System, Power State, Status, and Service.

Discovered Services

You can view discovered services, the number of VMs on which each discovered service is running, and you can configure service discovery.

Where You View the Discovered Services

From the menu, select **Home**, and then from the left pane select **Discover Services**.

Discovered Services

You see a list of services that are discovered and the number of VMs that have the services running. You see this section after you have configured Service Discovery and the services are discovered.

Known Services

You see a list of all the services supported and those that can be discovered.

Whitelisted Services

You can configure a service by clicking **Configure Whitelist**, and adding a process name, port, and display name in the **Whitelist Service** dialog box.

The process name must exactly match the name that you see in the guest OS when running commands `ps` in Linux and `wmic` in Windows. Specify a single port for each service.

Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

Table 1-55. Virtual Machine Metrics

Metric Name	Description
Guest OS Services|Total Number of Services	Number of out-of-the-box and user-defined services discovered in the VM.
Guest OS Services|Number of User Defined Services	Number of user-defined services discovered in the VM.
Guest OS Services|Number of OOTB Services	Number of out-of-the-box services discovered in the VM.
Guest OS Services|Number of Outgoing Connections	Number of outgoing connection counts from the discovered services.
Guest OS Services|Number of Incoming Connections	Number of incoming connection counts to the discovered services.

Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

Table 1-56. Service Summary Metrics

Metric Name	Description
Summary|Incoming Connections Count	Number of incoming connections.
Summary|Outgoing Connections Count	Number of outgoing connections.
Summary|Connections Count	Number of incoming and outgoing connections.
Summary|Pid	Process ID.

Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

Table 1-57. Service Performance Metrics

Metric Name	Description
Performance metrics group|CPU	CPU usage in percentage.
Performance metrics group|Memory	Memory usage in KB.

Table 1-57. Service Performance Metrics (continued)

Metric Name	Description
Performance metrics group|IO Read Throughput	IO read throughput in KBps.
Performance metrics group|IO Write Throughput	IO write throughput in KBps.

Service Type Metrics

Service discovery discovers metrics for service type objects.

Table 1-58. Service Type Metrics

Metric Name	Description
Number of instances	Number of instances of this service type.

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 search for log messages and collect and analyze log feeds. You can view log-related metrics for troubleshooting. 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.

You can integrate only one vRealize Log Insight instance.

Prerequisites

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

- Verify that you know the IP address or FQDN of the vRealize Log Insight instance you have installed.

Procedure

1   In the menu, select **Administration**, and then from the left pane, select **Management > Integrations**.

2   From the **Integrations** page, click VMware vRealize Log Insight.

3   In the VMware vRealize Log Insight page complete the following steps:

- Enter the IP address or FQDN in the **Log Insight server** text box of the vRealize Log Insight you have installed and want to integrate with.

- Select the collector group from the **Collectors/Groups** drop-down menu.

- Click **Test Connection** to verify that the connection is successful.

- Click **Save**.

4   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.

5   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 **vRealize Operations** from the left pane. You see the vRealize Operations Integration pane.

2   In the **Hostname** text box, enter the IP address or FQDN of the vRealize Operations Manager instance you want to integrate with.

> **Note**   If you are using a load balancer, use its IP address or FQDN as a hostname value.

3   In the **Username** and **Password** text boxes, enter the user name and password of the vRealize Operations Manager instance you want to integrate with.

4   Select the relevant check boxes according to your preference:

   ■   To send alerts to vRealize Operations Manager, select **Enable alerts integration**.

   ■   To let vRealize Operations Manager open Log Insight and query for object logs, select **Enable launch in context**.

   ■   To calculate and send metrics to vRealize Operations Manager, select **Enable metric calculation**.

5   Click **Test Connection** to verify that the connection is successful and accept the certificate if it is untrusted.

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 1-59. Log Forwarding Page Options

Options	Description
Self-monitoring logging configuration	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.
Protocol	You can select either `cfapi` or `syslog` from the drop-down menu to send event logging messages.
Port	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: <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>
Use SSL	Allows the vRealize Log Insight agent to send data securely.

Table 1-59. Log Forwarding Page Options (continued)

Options	Description
Path to Certificate Authority File	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 `/etc/pki/tls/certs/ca-bundle.crt` or `/etc/ssl/certs/ca-certificates.crt`.
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.

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	`Maximum([(10000-0)*0.4],[(10000-0)/5])`
				`= Maximum(4000, 2000) = 4000`
				
				which is 333.33 per month.

Year 2	10000	0.4	4000	`Maximum([(10000-4000)*0.4],[(10000-4000)/4])`
				`= Maximum (2400, 1500) = 2400`
				
				which is 200 per month.

Year 3	10000	0.4	6400	`Maximum([(10000-6400)*0.4],[(10000-6400)/3])`
				`= Maximum (1440, 1200) = 1440`
				
				which is 120 per month.

Year 4	10000	0.4	7840	`Maximum([(10000-7840)*0.4],[(10000-7840)/2])`
				`= Maximum (864, 1080) = 1080`
				
				which is 90 per month.

Year 5	10000	0.4	8920	`Maximum([(10000-8920)*0.4],[(10000-8920)/1])`
				`= Maximum (432, 1080) = 1080`
				
				which is 90 per month.

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.

You can now set a total cost for the License, Labor, Network, Maintenance, and facilities cost drivers in vRealize Operations Manager :

Note   The total cost set by you is distributed across resources in the data center. For example, if you set the total cost for the RHEL license, the cost is divided across all the hosts and VMs which use the RHEL license.

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 might 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 the reference database. Based on your inputs, vRealize Operations Manager recalculates the total amount for

the 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.

- Datastores: 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 the 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 invRealize 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 1-60. Expense Types

Cost Drivers	Description
Server Hardware : Traditional	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.
	Note You can now select an individual server from the server group and specify the unique cost for each individual server.
Server Hardware : Hyper-Converged	The Server Hardware : Hyper-Converged cost driver, tracks the expenses associated with hyper converged infrastructure components. The Server Hardware : Hyper-Converged cost driver includes expenses for the Hyper Converged servers like vSAN enabled servers and vXRail. The expense provided is for both compute and storage.
	Note The customizations that were performed for vSAN server costing under Server Hardware : Traditional in the earlier versions will not be carried forward to 7.5 as the vSAN enabled servers will fall under Server Hardware : Hyper-Converged servers now.
Storage	You can calculate the storage cost at the level of a datastore based on the tag category information collected from vCenter Server. You see the storage total distribution based on category and the uncategorized cost details.
	Note The vSAN datastores are not displayed as part of this cost driver page.

Table 1-60. Expense Types (continued)

Cost Drivers	Description
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 ■ 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 the 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.
Application Cost	You can see the cost of different application services you are running in your environment compared to your overall expenses. Some examples of application cost are, cost of running SQL server cluster and cost of running Antivirus on VMs.

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

Cloud Providers Overview

By default, you can see that Amazon Web Services (AWS), Google Cloud, IBM Cloud, and Microsoft Azure are included in vRealize Operations Manager . You can also add your own cloud provider by using a standard vRealize Operations Manager template.

You can configure the new cloud provider as per the standard vRealize Operations Manager template and perform a migration scenario. The vRealize Operations Manager template contains data points for vCPU, CPU, RAM, OS, region, plan term, location, and built-in instance storage, you must provide these values when you add cloud providers. The result of the migration scenario helps you assess the cost savings achieved using your cloud provider against the default cloud providers.

You can edit the rate card for new cloud providers and default cloud providers. However, you cannot delete the default cloud providers.

Add Cloud Provider

You can use the Add Cloud Provider workspace to add or edit a cloud provider. You can edit the cloud provider rate card for default cloud providers and the new cloud provider.

Procedure

1   On the menu, click **Administration** and in the left pane click **Configuration** > **Cost Settings** > **Cloud Providers**.

    You can also reach the Cloud Providers page from the Home Screen. In the Home screen, navigate to **Optimize Capacity** > **What-If Analysis** > **Plan Migration** > **Add Cloud Providers**. For more information, see **What-If-Analysis - Migration Planning** section in vRealize Operations Manager help.

2   Click the **Add Cloud** icon.

3   Enter the **Cloud Provider Name**.

4   Select the cloud provider logo and click **Upload Logo**.

5   Click **Next**.

6   Click **Download Template** and specify the required values.

    Note   When you edit a cloud provider the Download Template link is replaced with Download Existing Rate Card. You can update the existing rate card and upload the same.

7   Select the updated template and click **Upload Rate Card**.

8   Click **Validate**.

    Note   vRealize Operations Manager validates the rate card and reports success or failure. If errors are reported, you can correct the errors and proceed further.

9   Click **Finish**.

Results

The new cloud provider is now part of the vRealize Operations Manager cloud provider list.

Billing Framework for Unmanaged Objects

You can remove objects which should not be monitored by vRealize Operations Manager using the billing framework. The billing framework ensures that the license fee is not applicable to the unmanaged objects which are moved to the maintenance state.

How To Manage Unmanaged Objects

To manage the unmanaged objects, you have to perform the following actions in vRealize Operations Manager :

- Remove objects that should not be monitored.

- Move the unmonitored objects to maintenance state.

- Stop the data collection for the objects in maintenance mode.

- Power off the virtual machines that are in maintenance mode.

Billing Support for Unmanaged Objects

When you remove specific objects from monitoring, vRealize Operations Manager moves these objects to maintenance mode and stops billing for the objects. The billing framework ensures that the costs related to licensing are not calculated for the following scenarios:

- vSphere and Public cloud virtual machines are in maintenance mode.

- vSphere and Public cloud virtual machines are in powered off state.

- vSphere and Public cloud virtual machines have stopped data collection.

The licensing fee is not charged for the objects in maintenance mode, you can verify the same in the next hourly billing cycle. You can navigate to **Administrator** > **Inventory** list, to view the list of objects that are in maintenance mode.

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 : Traditional

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. You can also specify the server cost for individual servers in a server group. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

Procedure

1   Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2   In the Cost Drivers tab, click **Server Hardware : Traditional**.

   Note   You can customize the default value of cost per server and specify exclusive values for other servers in the list.

   For example, if you have a system that has eight servers you can modify the default reference value from $1000 to $800 for eight servers. You can also select two servers from the list and customize their value as $600. So, any new server that is added to the system will have the default value as $800.

3   Select the required edit mode for changing the server hardware cost.

   - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 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.

5 After selecting a server group, you can manually enter the required fields.

a Enter the Purchase Type and Cost Per Server.

  Note You can use the **+ ADD COST PER SERVER** option to create multiple server batches and set the cost for a specific server in a server group.

b Click **Save**.

Editing Server Hardware: Hyper-Converged

You can view, add, edit, or delete the cost of Hyper converged Infrastructure (HCI) component in your server group. You can specify the cost per server and compute percentage exclusively for the HCI servers. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

Procedure

1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2 In the Cost Drivers tab, click **Server Hardware : Hyper-Converged**.

3 Select the required edit mode for changing the server hardware cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

■ **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 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 servers falling under vSAN clusters and vXrail servers 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.

Note You can edit the Compute Pct column to adjust the storage rate of the vSAN datastores. You can use the same percentage to determine the cost.

5 After selecting a server group, you can manually enter the required fields.

a Enter Purchase Type, Cost Per Server, and Compute Percentage.

Note You can use the **+ ADD COST PER SERVER** option to create multiple server batches and to customize the cost per server.

b 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 the 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
Edit Mode	You can select the storage cost to be applicable for all the data centers or a specific data center. ■ **Edit for All Data Centers** mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost. ■ **Edit for specific Data Center** mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.
Select Data center	You can select the data center for which you want to change the storage cost. This field is applicable only for specific data centers.
Tag Category	■ **Category** displays the tag categories for datastores and also the tags associated with the category. **Note** If you perform fresh installation of vCenter Server 6.0, and not assign tags to the datastores, cost drivers display the tag category for datastores as `uncategorized`.
Datastores	Displays the total number of datastores for a specific category or type. You can click the datastore value to see the 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 now set a total fixed cost for the license in vRealize Operations Manager . The total license cost is divided across all the hosts present in the data center. 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**.

3 Select the required edit mode for changing the license cost.

■ **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 Click **Save**.

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**.
	Note Two new cost components, Monthly cost of VMware vSAN Per Socket and Monthly cost of VMware vSAN SnS have been included for the vSAN cost calculation. The default values for these components are based on the reference database values.
	The licensing cost for the Windows operating system falls under one of the following categories:
	Per Core License, applicable for
	■ Windows Server 2016
	■ Windows Server 2019
	Per Scoket License, applicable for
	■ Windows NT 4.0
	■ Windows Server 2003
	■ Windows Server 2008
	■ Windows Server 2012
	Per Instance License, applicable for
	■ Windows XP
	■ Windows Vista
	■ Windows 98
	■ Windows 95
	■ Windows 8
	■ Windows 7
	■ Windows 3.1
	■ Windows 2000
	■ Windows 10
VMs	Displays the number of virtual machines that are running on the specific operating system.
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.
	Note The Charged By column can be edited to mention that the cost is charged by socket, core, instance, or ELA.
Total Cost	Displays the total cost of the specific operating system.

5 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.

Customizing License Assignment

You can customize the licensing cost associated with your host using the custom license assignment option. Based on your requirement you can add or delete different operating system licenses to your host. With the custom license assignment option, you can increase or decrease the licensing cost associated with your host.

Procedure

1 Click **Administration** and in the left pane click **Configuration > Cost Drivers**.

2 In the Cost Drivers tab, click **License**.

3 Select the required edit mode for changing the monthly license cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 To customize the license cost for a specific server, click **Customize License Assignment**.

5 Select the host for which you want to customize the license cost and click **Assign**.

6 From the drop-down menu, select the operation system and click **Ok**.

The new operating system is listed under the Current Assignment column.

7 To remove an existing operating system from the host, under **Current Assignment** click X icon next to the operating system.

The license cost of the removed operating system is reduced from the total cost.

8 Click **Save**.

9 Navigate to the **Cost Calculation Status** tab and click **Run**.

Results

The license cost is updated for the host, the * sign next to the host indicates that the license cost for the host has changed.

Category	Description
Server	You can select the server for which you want to customize the license cost.
Current Assignment	Displays the current operating systems associated with the host.
Default Assignment	Displays the default operating systems associated with the host.
Filter	Filters the hosts based on the operating system type.
Reset	Resets the license cost of the host to the default value.

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. You can snow specify a total fixed cost for maintenance in vRealize Operations Manager . The total maintenance cost is divided across all the hosts present in the data center.

Procedure

1   Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2   In the Cost Drivers tab, click **Maintenance**.

3   Select the required edit mode for changing the monthly maintenance cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  Note   When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4   To customize the maintenance cost for a specific server, click **Edit For Individual Servers**.

5   Click **+Add Cost Per Server**.

6   From the **Select Server's for customization** drop-down select the required server and click **Ok**.

7   Specify the Server Hardware Percentage and OS Percentage and click **Save**.

View the change in maintenance cost after you have run the cost calculation cycle.

Edit Monthly Cost of Labor

You can edit the monthly cost of labor for your cloud environment. You can set a total fixed cost for labor in vRealize Operations Manager . The total labor cost is divided across all the hosts present in the data center. 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**.

3 Select the required edit mode for changing the monthly labor cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

   Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 Edit the monthly labor cost.

- Edit the detailed cost of labor.

- Edit the total monthly labor cost for servers, virtual infrastructure, and operating system.

5 To customize the labor cost for a specific server, click **Server** and then click **Edit For Individual Servers**.

6 Click **+Add Cost Per Server**.

7 From the **Select Server's for customization** drop-down select the required server and click **Ok**.

8 Specify the Monthly hours of labor per hour, Labor hourly rate, and click **Save**.

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

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. You can now set a total fixed cost for network resources in vRealize Operations Manager . The total network cost is divided across all the hosts present in the data center.

Procedure

1   Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2   In the Cost Driver tab, click **Network**.

3   Select the required edit mode for changing the monthly network cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

    Note   When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4   Edit the monthly cost of network.

- Modify the values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and the 100 Gigabit NIC.

- Modify the total monthly cost of all network expenses associated with the cloud.

5   To customize the network cost for a specific server, click **Edit For Individual Servers**.

6   Click **+Add Cost Per Server**.

7   From the **Select Server's for customization** drop-down select the required server and click **Ok**.

8   Specify values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and 100 Gigabit NIC and click **Save**.

    View the change in network cost after you have run the cost calculation cycle.

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. You can now set the total fixed cost for facilities in vRealize Operations Manager . The total facilities cost is divided across all the hosts present in the data center.

Procedure

1   Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2   In the Cost Driver tab, click **Facilities**.

3   Select the required edit mode for changing the monthly facilities cost.

   ▪   **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

   ▪   **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

4   (Optional) Select the data center from the drop-down menu.

   Note   If you select Edit for specific data center as the edit mode, then the select data center option is enabled.

5   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.

6   To customize the facilities cost for a specific server, click **Edit For Individual Servers**.

7   Click **+Add Cost Per Server**.

8   From the **Select Server's for customization** drop-down select the required server and click **Ok**.

9   Specify the Cost Per Kilowatt and Real Estate Cost Per Rack Unit and click **Save**.

   View the change in network cost after you have run the cost calculation cycle.

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   Additional cost driver allows you to assign costs at Host, vCenter, VM, cluster, or data center level. For example, if you want to keep a cluster protected using the disaster recovery services, which involves an additional cost of $5000, you can do that by editing the additional cost driver.

4   Select the **Entity Type** and **Entity Selection**.

The **Entity Count** gets updated.

5   Enter the **Monthly Cost per entity** .

The **Total Cost per month** gets computed automatically.

6   Click **Save**.

Note   After you update the Additional Cost configuration, you must reload the page manually to view the updated values.

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.

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   Click the edit icon next to the application cost you want to edit.

Note   You can now specify the cost of packaged applications that are discovered by the Service Discovery Management Pack. Earlier the option to specify the application cost was available only for business applications defined by the user.

4   Modify the cost of the application.

5   Click **Save**.

Cluster Cost Overview

vRealize Operations Manager calculates the base rates of CPU and memory so that they can be used for the 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 a 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. Note When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value.
Expected Memory Utilization	The utilization percentage level of Memory that the cluster is expected to operate. Note When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value.

Cluster Cost Computation with Allocation Model

You can now use the allocation model to compute the cost of clusters in vRealize Operations Manager , earlier the cluster cost computation was based on the cluster utilization. When you perform cost computation using the allocation model, you can set the over commit ratio for CPU, RAM, and storage.

Note The allocation ratio can be set at both cluster level and datastore cluster level. You can also mention the storage base rate, which will displayed at the datastore level.

Table 1-61. Cluster Base Rate Computation with Allocation Model

Base Rate	Formula
vCPU Base Rate	vCPU base rate = B1 = (Cost attributed to CPU) / (Number of vCPUs in a cluster)
RAM Base Rate	RAM base rate = B2 = (Cost attributed to RAM) / Number of vRAMs in a cluster)
	Note The cost computation is based on Over Commit ratio. If the Over Commit ratio is 1:4, and total cores in cluster are 6, then vCPU count = 24, in case if the allocated vCPU exceeds this targeted number, then the maximum value is selected.

Table 1-62. Virtual Machine Cost Computation with Allocation Model

Cost	Formula
Virtual Machine Cost	Virtual machine cost = (Number of vCPU allocated x B1 of cluster it belongs to) + Number of vRAMs allocated x B2 of cluster it belongs to) + storage cost + direct cost.
	Note Storage allocated represents the Storage Base Rate based on allocation.

Editing Cluster Cost Calculation Methods

You can edit the cluster cost calculation method based on your business requirement. The cost of a cluster is derived from cost drivers. Virtual machine cost is calculated by multiplying base rates with the utilization of the VMs.

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
Cluster Usable Capacity After HA and Buffer	The cluster cost calculated total capacity minus resources needed for High Availability (HA) and the capacity buffer setting.
	Base rates are calculated based on the total cost of the cluster and Usable Capacity after HA and Buffer. Virtual machine costs are calculated from these base rates. Things to note: ■ A lower buffer reduces the base rates and causes the virtual machines to become cheaper. ■ A higher buffer increases base rates and causes the virtual machines to become more expensive. ■ Base rates and virtual machine costs do not change with the utilization of the cluster. ■ The difference between Usable Capacity after HA and Buffer and actual utilization is used to compute unallocated costs.
Cluster Actual Utilization	To calculate the base rates using the month to date average utilization of the cluster resources, select this option.
	Base rates are calculated based on the total cost of the cluster and average utilization. Virtual machine costs are calculated from these base rates. Things to note: ■ Lower utilization level causes base rates to be high and virtual machines also become more expensive. ■ Higher utilization level causes base rates to be lower and virtual machines to become cheaper. ■ Base rates and virtual machine costs can change frequently based on the utilization of the cluster. ■ Unallocated cost of the cluster is near to zero. ■ The costs for unused resources are distributed across all virtual machines based on their actual utilization within the cluster.

4 Click **SAVE**.

Publish Daily Cost Metrics for Virtual Machines

In vRealize Operations Manager , you can now publish daily cost metrics for all virtual machines. The daily cost metric of a virtual machine is the sum of daily cost of CPU, memory, storage, and additional cost associated with the virtual machine. Daily cost metrics provide granular details of the costs associated with the virtual machine.

Formula to Calculate the Daily Cost and Monthly Cost of Virtual Machines

You can calculate the daily cost associated with a virtual machine using the following formula.

Virtual Machine Cost Elements	Calculation
Daily Total Cost of Virtual Machine	Daily total cost of virtual machine = Sum of Daily cost of (CPU + memory + storage + additional cost)

The change in daily cost metrics also changes the way you calculate the effective month to date cost of a virtual machine. You can use the following formula to calculate Effective Month to Date cost for a virtual machine.

Virtual Machine Cost Elements for a Month	Calculation
Effective MTD Cost of VM	Sum of CPU daily cost from the beginning of the month until now + Sum of memory daily cost from the beginning of the month until now + Sum of storage daily cost from the beginning of the month until now + Sum of additional daily cost from the beginning of the month until now

How to View the Daily Cost Metrics of a Virtual Machine

To view the daily cost metrics of a virtual machine, from the menu, select **Administrator** and then in the left pane select **Inventory** > **vCenter Adapter** , select the specific **Virtual Machine**, and click the **Metrics** tab.

Pricing Overview

You can create pricing cards in vRealize Operations Manager to calculate the price associated with your virtual infrastructure. You can assign pricing cards to vCenters or Clusters, depending on the pricing strategy determined by vRealize Operations Manager administrator. The pricing cards help you to set the price for each resource present in your virtual environment.

You can customize the pricing card as per your requirement. vROps has two types of pricing cards, rate-based pricing card and cost-based pricing card. After configuring a pricing card, you can assign it to one or more vCenters or Clusters as determined by the pricing strategy.

How Is Price Calculated

In rate-based pricing policy vRealize Operations Manager calculates the virtual infrastructure price based on the rate card defined by you. For rate-based pricing policy vRealize Operations Manager lets you define cost elements as per your requirements.

The server recalculates the price every 24-hours, the price calculation for the new pricing cards is done in the next vRealize Operations Manager price calculation cycle.

Hierarchy of Pricing Policy

The assignment of policy in vRealize Operations Manager will be for Clusters and vCenters. The price is calculated for virtual machines, then it is aggregated and rolled up to vCenter. If there are two policies, a default policy for vCenter and another policy for Cluster, then the price calculation is based on the cluster policy for all the resources under the cluster. After that the cluster cost is rolled up to vCenter.

When a virtual machine is under vRealize Automation hierarchy and vCenter hierarchy, then the pricing is calculated based on the vRealize Automation hierarchy and the virtual machine is removed from the vCenter resources and included under vRealize Automation resources.

Pricing Support for VMware Cloud on AWS Resources

You can create a pricing policy in vRealize Operations Manager and assign it to VMware Cloud on AWS (VMC) resources, however you can only use the rate-based pricing policy for VMC-related objects.

Note When you assign Cost-Based Policy for VMC resources, the policy is not applied, and price calculated for the policy is reported as zero.

Add New Pricing Card

You can add and assign new pricing card to vCenter and Clusters in vRealize Operations Manager . The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

Procedure

1 Navigate to **Administration** > **Cost Settings** > **Pricing**.

2   Click **New Pricing Card** and configure the details of the pricing card.

Table 1-63. Pricing Card Configuration

Parameter	Description
Name and Description	1 Enter a name and description for your pricing card. 2 Optional: Select Default for Unassigned Workloads. 3 Click **Next**. Default pricing card applies to all vCenter resources which do not have a direct cost policy assigned to them.

Basic Charges	Select the type of pricing card. Follow the steps for cost-based pricing card.
	1 Enter the Cost Factor for the following.
	   a CPU Cost
	   b Memory Cost
	   c Storage Cost
	   d Additional Cost
	2 Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**.
	3 Select how to charge for the resources, the options are **Always** or **Only When Powered On**.
	4 Click **Next**.
	
	Note Cost - The cost is defined in vRealize Operations. If selected, a multiplication factor is required. For example, if you select 1.1 as a factor, the cost is multiplied by 1.1 resulting in a 10% increase to the calculated cost. The price equation using cost is: <cost> x <multiplication factor> = Price
	
	Follow the steps for rate-based pricing card.
	1 Enter the CPU Rate in MHz per vCPU.
	2 Enter the Memory Rate per GB.
	3 Enter Storate Rate per GB.
	4 Select the ChargingPeriod for all the values.
	5 Select the Charge On Power State for all the values.

Guest OSes	1 Enter the Guest OS Name. 2 Enter the base rate. 3 Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**.

Table 1-63. Pricing Card Configuration (continued)

Parameter	Description
Tags	Enter the Tag name and Tag Value. Define the charging method and base rate. ■ Recurring - enter a base rate and define recurring interval as the charge period. The absolute rate value is required and it is added to the overall price. ■ One time - define the one-time base rate charge. The absolute value is required and it is added as a one time price. ■ Rate Factor - A multiplication factor is required that is applied to the select charge category. Select how to charge the Tag based on powered on state.
Overall Charges	You can define overall charges to VMs that match this policy. 1 Enter the VM setup charges. 2 Enter the Recurring charge and select the time period from the drop-down menu.
Assignments	You can assign the new pricing card to vCenters and Clusters. 1 Select the vCenter or Cluster to which you want to apply the pricing card. 2 Click **Add** and Click **Finish**.

Results

The new pricing card details are displayed in the Pricing tab.

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 7.5.

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

Costing Enhancements

In vRealize Operations Manager , a new global property Cluster Utilization Ceiling Factor is introduced. Using Cluster Utilization Ceiling Factor, you can specify the ceiling value and calculate the base rate for a cluster.

You can use the ceiling factor only if the base rate cost calculation is done using Cluster Actual Utilization method. After you set the ceiling factor value, the Actual Utilization of the cluster is rounded off to the next available multiple of the ceiling value. When ceiling value is 0, Expected Utilization is equal to actual utilization. When ceiling value is 20, it is not considered as special case, actual utilization is rounded off to the next multiple.

Note   The ceiling value range is from 0 to 20. If the number is out of this range, the default value of five is used as the ceiling number.

How to Set the Cluster Base Rate Calculation Method

To change the Cluster Base Rate Calculation method, you must go to **Administration** > **Configuration** > **Cost Settings** > **Cluster Cost** page. Click **Change** next to the Cluster Base Rate calculation method and select Cluster Actual Utilization.

Where to Find Cluster Utilization Ceiling Factor

To set the ceiling value for a cluster, you must go to **Administration** > **Management** > **Global Settings** > **Cluster Utilization Ceiling Factor**. Enter the ceiling value between 0 and 20 and click **Save**.

To view the change is cost metrics, run the Cost Calculation Status and select a cluster .

If the Actual Utilization of the cluster for CPU is 30 % and Memory is 45%, and the ceiling value specified is 10, then

- Cluster Expected CPU Utilization (%) = 40

- Cluster Memory Expected Utilization (%) = 50

Actual Cluster Utilization is rounded off to the ceiling value.

If you set the Cluster Utilization Ceiling Factor to either 0 or 20, then the value of Expected Memory Utilization changes to the next number. For example, if you set the ceiling factor to 0 then, the expected utilization value changes to 1.

Support to Roll up Name Space Cost Metrics

The cost metrics of Point of Delivery (Pod) virtual machines (VMs) has been enhanced to support the following scenarios:

- Cost metrics of Pod VMs are rolled up to the Name Space and Guest Cluster level.

- All the cost metrics of VMs, Pods, and guest cluster which are present under Name Space are rolled up to Name Space and Guest Cluster level.

Old Cost Metrics	Rolled up Cost Metrics
Effective MTD Total Cost	Aggregate Additional Daily Cost
Deleted VM Daily Cost	Aggregate Deleted VM Daily Cost
Daily CPU Cost	Aggregate CPU Daily Cost
Daily Memory Cost	Aggregate Memory Daily Cost
Daily Storage Cost	Aggregate Storage Daily Cost
Daily Additional Cost	Aggregate Additiona Daily Cost

vRealize Automation 7.x

vRealize Automation 7.x extends operational management capabilities of the vRealize Operations Manager platform to provide tenant-aware operational visibility of the cloud infrastructure.

vRealize Automation 7.x 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 vRealize Automation 7.x 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. vRealize Automation 7.x 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.

- To view capacity for tenants, business groups, and reservations. From the menu, select **Administration** and then in the left pane, select **Inventory**. Select the **Objects** tab in the right pane. By default, the usage capacity model is enabled for these objects. You can enable the allocation model from the policy settings.

Supported vRealize Automation Versions

vRealize Operations Manager 8.x 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 8.2, vRealize Automation Management Pack is upgraded to 8.0.

Object Types and Relationships

vRealize Automation 7.x 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

- User

  You can view the different users from the **Inventory > List** tab. The user object type has a relationship with VMs, deployments, and business groups.

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 1-64. Relationship Model

Relationship View	Parent-Child Relationship Between Objects
vRealize Automation Tenant View	Tenant > Business Group > Reservation
vRealize Automation App View	Tenant > Blueprint > Deployment > VM
vRealize Automation Custom Data Center View	CDC > Cluster > Host > VM
vRealize Automation Reservation Policy View	Reservation Policy > Reservation > VM
vRealize Automation Virtual Machine View	Tenant > Business group > Deployment > VM

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 vRealize Automation 7.x 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   vRealize Automation 7.x supports only vCenter objects used and managed by vRealize Automation. No other object kinds such as AWS or Openstack resources are supported currently.

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 can 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.

vRealize Automation 7.x 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 are associated with that adapter instance. Any communication to the vRealize Automation servers ensures that the certificates presented by the servers match the ones accepted by the user.

Configuring vRealize Automation

You can configure an instance of 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 vRealize Automation 7.x 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, select **Administration**, and then from the left pane, select **Solutions > Repository**.

2 From the **Repository** page, on the right side, select VMware vRealize Automation 7.x from the Native Management Packs section, and click **Activate**.

vRealize Automation 7.x is installed and appears in **Other Accounts > Add Accounts** pane.

3 In the menu, click **Administration**, and then from the left pane click **Solutions > Other Accounts > Add Accounts**.

4 Click vRealize Automation 7.x and configure the solution.

Option	Description
Name	The name for the adapter instance.
Description	(Optional) The description of the adapter instance.
vRealize Automation Appliance URL	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`. If there is a load balancer for the CAFÉ appliances, the URL must have host name or IP address of the load balancer in the format `https://HostName` or `https://IP`.
Credential	To add the credentials to access the vRealize Automation environment, click the plus sign. ■ **Credential name**. The name by which you are identifying the configured credentials. ■ **SysAdmin Username**. The user name of the vRealize Automation system administrator. For information about the System Administrator, see System-Wide Role Overview. ■ **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	The collector on which the vRealize Automation 7.x runs. ■ 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.

Option	Description
Tenants	Collects data for specific tenants associated with vRealize Automation. To collect data, configure the tenants in the following manner: ■ * (by default). Data is collected for all tenants. > **Note** > ■ 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 vRealize Automation 7.x 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. ■ **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	■ **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.
vRealize Automation VA FQDN	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. Enter these details only when you enable vRealize Automation system health monitoring.
vRealize Automation adapter collection interval (minutes)	The time interval between data collections by vRealize Automation 7.x. 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. To change this value to less than 5 minutes, you must change the collection interval value in the adapter.
Tenant resource collection interval (minutes)	The time interval between the data collected by the tenants in vRealize Automation 7.x. 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. To change this value to less than 5 minutes, you must change the collection interval value in the adapter.

Option	Description
Business group resource collection interval (minutes)	The time interval between the data collected by the business groups in vRealize Automation 7.x.
	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.
	To change this value to less than 5 minutes, you must change the collection interval value in the adapter.
Blueprint resource collection interval (minutes)	The time interval between the data collected by the blueprints in the vRealize Automation 7.x.
	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.
	To change this value to less than 5 minutes, you must change the collection interval value in the adapter.
Autodiscovery	Discover objects automatically.
	■ To set automatic discovery for objects, select **True**.
	■ To set off the automatic discovery, select **False**.

5 Click **Test Connection** to validate the connection.

If one of the tenant connections is successful, Test Connection is successful.

6 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 1-65. 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

Table 1-65. Configuration Properties (continued)

Property Name	Description	Default Value
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 **Note** It is recommended that you keep the maximum value as 100. Refer to the sizing guidelines for large-scale environment guidelines: Sizing Guidelines

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.

vRealize Automation 8.X

The vRealize Automation 8.x extends operational management capabilities of the vRealize Operations Manager platform to provide the cloud aware operational visibility of the cloud infrastructure. The vRealize Automation 8.x enables you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use the vRealize Automation 8.x to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with vRealize Operations Manager .

- Import and synchronize existing cloud accounts from vRealize Automation 8.x to vRealize Operations Manager .

- Manage the workload placement of VMs that are part of the clusters managed by vRealize Automation 8.x.

- Integrate and troubleshoot vSphere endpoint issues associated with vRealize Automation 8.x using the vRealize Operations Manager dashboard.

  Note   In this release we support only vSphere endpoints.

vRealize Operations Manager and vRealize Automation Integration - Technical Overview

The vRealize Automation 8.x integration with vRealize Operations Manager , extends operational management capabilities of the vRealize Operations Manager platform to provide cloud aware operational visibility of the cloud infrastructure. The vRealize Automation 8.x enables you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use the vRealize Automation 8.x to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with vRealize Operations Manager .

- Import and synchronize existing cloud accounts from vRealize Automation 8.x to vRealize Operations Manager .

- Manage the workload placement of VMs that are part of the clusters managed by vRealize Automation 8.x.

- Integrate and troubleshoot vSphere endpoint issues associated with vRealize Automation 8.x using the vRealize Operations Manager dashboard.

How Does vRealize Automation and vRealize Operations Manager Integration Work

vRealize Automation can work with vRealize Operations Manager to perform advanced workload placement, provide deployment health and virtual machine metrics, and display pricing.

Integration between the two products must be on-premises to on-premises, not a mix of on-premises and cloud.

To integrate with vRealize Operations Manager , look under **Infrastructure** > **Connections** > **Integrations**. To add the integration, you need thevRealize Operations Manager URL and its login user name and password. In addition, vRealize Automation and vRealize Operations Manager need to manage the same endpoint.

Workload Placement

When you deploy a blueprint, workload placement uses collected data to recommend where to deploy the blueprint based on available resources. vRealize Automation and vRealize Operations Manager work together to provide placement recommendations for workloads in the deployment of new blueprints.

While vRealize Automation manages organizational policies, such as business groups, reservations, and quotas, it integrates with the capacity analytics of vRealize Operations Manager to place machines. Workload placement is only available for vSphere endpoints.

Workload Placement Terms Used

Several terms are used with workload placement.

- Clusters in vSphere map to compute resources in vRealize Automation.

- Reservations include compute and storage, where the storage can consist of individual datastores or datastore clusters. A reservation can include multiple datastores, datastore clusters, or both.

- Multiple reservations can refer to the same cluster.

- Virtual machines can move to multiple clusters.

- When workload placement is enabled, the provisioning workflow uses the placement policy to recommend where to deploy the blueprint.

Provisioning Blueprints with Workload Placement

When you use workload placement to provision blueprints, the provisioning workflow uses the reservations in vRealize Automation, and the placement optimization from vRealize Operations Manager .

1   vRealize Operations Manager provides placement optimization recommendations according to analytics data.

2   vRealize Automation continues the provisioning process according to the placement recommendations from vRealize Operations Manager .

If vRealize Operations Manager cannot provide a recommendation, or the recommendation cannot be used, then vRealize Automation falls back to its default placement logic.

Workload Placement Goals

The goal of Workload Placement is to make sure that no cluster is overloaded by more than 80% of the potential workload. Workload placement is done in the following three stages.

Stress-free clusters

Ensures that the memory, CPU, or disk space workload is less than 80% for the cluster.

Workload Placement based on Business Intent

Distribution of virtual machines between Clusters is based on tags. When a cluster and VM have the same tag, VM will be recommended to move from this cluster or VM will be recommended to move to this cluster. When host-based tagging is enabled, VM will be recommended to optimize the workload for cluster based on a rule.

Distribution Strategy

- Balanced distribution: Distribution is based on the green zone, with maximum 20% difference workload between the two clusters.

- Moderate distribution: Ensures that no cluster is at stress level.

- Consolidated distribution: Keeps the hosts free while maintaining the workload at green level. In some cases, one of the clusters has resources free for backup purposes.

Workload Placement Recommendation

Workload Placement is recommended to run on a cluster (with existing VMs) or for the new deployment in vRealize Automation for Day 0 integration. If after the deployment or move of the virtual machine, the cluster hosting that VM does not have workload greater that 80% for CPU and/or Memory and/or Disk Space. The recommendation starts only if the Memory or CPU workload is not optimized.

Note   We do not recommend Disk Space Optimization for Workload Placement, as we always ensure that the workload for Disk Space is within the green zone.

vRealize Automation Workload Placement Day 1 Recommendation

The distribution of VMs is done based on blueprint configurations. WLP calculates and evaluates the impact of potential deployment based on the workload or cluster utilization. The objective of WLP is to make sure that the least loaded cluster gets to provision highest number of VMs.

We have a cluster which has 100 GB memory capacity of which 20 GB is free, so this means 80% free space. We have another cluster which has 1 TB memory of that 300 GB is used. If you look at this percentage wise, we have 70% free space, however in terms of actual available space, we have 700 GB free space which is more than the 20 GB free space available in the first scenario.

Note   If the workload placement results in having more than 80% workload on the cluster, thenvRealize Operations Manager cannot provide a recommendation, or the recommendation cannot be used, then vRealize Automation falls back to its default placement logic.

Workload Placement Automation

Automation

The Automation calculates and evaluates the move of virtual machines every 5 minutes. If you find a VM that is not optimized, the optimization is triggered automatically. Note, the time slot between two automated optimizations is limited to 6 hours.

Schedule

Schedule automation calculates and evaluates the move only during the scheduled time slots. The available options are Once, Daily, Weekly, and Monthly.

Impact on Cloud Zones and Non vRealize Automation-Managed VMs

Whenever there is a vRealize Operations and vRealize Automation integration for a data center. The cloud zones that have virtual machines which are not managed by vRealize Automation or not created by vRealize Automation, Workload Placement ignores them.

Supported vRealize Automation Versions

vRealize Automation 8.x is supported on vRealize Operations Manager 8.2 version. 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.

Object Types

vRealize Automation 8.x brings in cloud accounts 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 .

- Cloud Zone

- Blueprint

- Project

- Deployment

- Cloud Account

- User

- Organization

- Cloud Automation Services World

Workload Placement

In vRealize Operations Manager , you can configure vRealize Automation 8.x instances to work with vRealize Operations Manager instances. Using vRealize Operations Manager you can monitor the placement of existing workloads and optimize the resource usage.

Prerequisites

- Verify that the user has privileges of Organizational Owner and Cloud Assembly Administrator set in vRealize Automation.

- You must know the vCenter Server credentials and have the necessary permissions to connect and collect data.

- Verify that vRealize Automation 8.x is enabled from **Administration** > **Management** > **Integrations** in vRealize Operations Manager . For more information, see Configuring VMware vRealize Automation 8.x with vRealize Operations Manager.

- vRealize Operations Manager must have the same vCenter Cloud Account configured to match with vRealize Automation 8.x.

- Ensure that integration is enabled for vRealize Operations Manager and vRealize Automation 8.x.

Procedure

1   In the menu, select **Home** and then select **Workload Optimization**.

2   Click the **View** filter drop-down menu and select the **VRA Managed** objects.

    All the Cloud Zones related to the vCenter Server are displayed in vRealize Operations Manager .

3   Click the **Cloud Zone** you want to optimize.

4   Based on the operational intent, click **Optimize Now**.

    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**.

6   Review the optimization moves, then click **BEGIN ACTION**.

    In the scope of vRealize Automation 8.x integration, vRealize Operations Manager sends a move migration request directly to vRealize Automation 8.x. In the earlier versions, the migration request was sent to the vCenter Server.

What to do next

To verify that the optimization action is complete, select **Administration** on the top menu, and click **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.

Pricing for vRealize Automation 8.x Components in vRealize Operations Manager

After you integrate vRealize Automation 8.x private cloud adapter instances with vRealize Operations Manager , you can calculate the cost of deployments, projects, and virtual machines of the selected cloud adapter. Pricing provides an overview of the costs related to the cloud environment, cloud resources, and the costs associated with the project.

How the Pricing Works in vRealize Automation 8.x

■ vRealize Operations Manager understands the constructs defined in vRealize Automation 8.x and calculates the CPU, RAM, Storage and Additional prices for Projects, Deployments, and virtual machines.

■ A single project can have multiple deployments and a single deployment can have multiple virtual machines associated with the deployment.

■ Pricing for multiple virtual machines associated with the deployment is the sum of all the resources associated with individual virtual machines.

■ If a single project has multiple deployments, then the project pricing is equal to the sum of individual deployments. The deployment can have multiple virtual machines and resources associated with it.

■ On day one, the pricing is equal to the cost of resources defined in vRealize Operations Manager .

■ On day two, the price is calculated using the following formula.

  ■ Cost of resources for the present day – Cost of resources for the previous day

■ If in case the pricing does not happen as per the definition, then the partial price is set to true, and the pricing is calculated based on the previous days price.

■ In vRealize Operations Manager , the following new dashboards are included to view the pricing details for the vRealize Automation 8.x instances.

  ■ Cloud Automation Environment Overview

  ■ Cloud Automation Project Cost Overview

  ■ Cloud Automation Resource Consumption Overview

  ■ Cloud Automation Top-N Dashboard

Data Collection Enhancements in vRealize Automation for Pricing in vRealize Operations Manager

The following enhancements have been made for the data collection process from vRealize Automation for pricing purposes.

■ Collect cloud zones with relation to clusters and resources pools from vRealize Automation to vRealize Operations Manager .

■ Collect Projects from vRealize Automation with relation to deployments.

■ Include project, cloud zone, and blueprint as properties in virtual machines that are deployed in vRealize Automation.

Upfront Price Support for vRealize Automation 8.x Private Cloud Components

vRealize Operations Manager supports upfront pricing for vRealize Automation 8.x in the following ways:

- vRealize Operations Manager uses rate cards to provide upfront cost estimates of catalog items just before deployment.

- vRealize Automation 8.x retrieves the deployment cost and estimated cost from vRealize Operations Manager .

- vRealize Automation user interface allows you to customize the pricing policies and assign them to the projects or cloud zones.

- If vRealize Automation does not specify the pricing policy, then the price is calculated using the vRealize Operations Manager cost calculation policy.

- If a custom pricing policy is set for a price calculation, then the deployment and upfront catalog price computation is done as per the custom policy.

Upfront Price Support for VMware Cloud on AWS Resources

vRealize Operations Manager supports upfront pricing for VMware Cloud on AWS resources in the following ways:

- vRealize Operations Manager supports upfront pricing for VMware Cloud on AWS only if rate-based pricing is configured in vRealize Automation for VMware Cloud on AWS resources.

- vRealize Operations Manager does not support cost-based computation for VMware Cloud on AWS resources.

Configuring VMware vRealize Automation 8.x with vRealize Operations Manager

To access vRealize Automation 8.x instance and troubleshoot automation issues using vRealize Operations Manager , you must configure the vRealize Automation adapter in vRealize Operations Manager .

Prerequisites

- Verify that you know the FQDN/IP address, user name, and password of the vRealize Automation instance you have installed.

- Ensure that the vRealize Automation user has both organizational owner and Cloud Automation Services administrator permissions.

Procedure

1 In the menu, select **Administration** and then from the left pane, select **Management > Integrations**.

2 From the **Integrations** page, click vRealize Automation8.x.

3   In the **vRealize Automation 8.x** page, enter the FQDN or IP address of the vRealize Automation 8.x instance to which you want to connect.

4   To add credentials, click the plus sign.

    a   In the Credential name text box, enter the name by which you are identifying the configured credentials.

    b   Enter the user name and password for the VMware vRealize Automation instance.

    c   Click **OK**.

    You have configured credentials to connect to a VMware vRealize Automation instance.

5   From the Collectors/Groups drop-down menu, select the collector group.

6   Click **Test Connection** to verify that the connection is successful.

7   Review and accept the Server certificate.

8   Click **Advanced Settings** and set Auto Discovery to true.

9   Click **Add** to save the adapter instance.

Results

After integrating vRealize Automation adapter instance with vRealize Operations Manager , you can view the vRealize Automation adapter data from the vRealize Operations Manager dashboard.

Cloud Zones in vRealize Operations Manager

Cloud zones enable you to group a set of compute resources and assign capability tags to the zone. The cloud zone is based on accounts/regions, so you must have at least one cloud account configured before you can create a cloud zone. Cloud zones define where and how blueprints configure deployments. You can have one or many cloud zones assigned to each project based on priority and limits.

How Cloud Zones Work

After you integrate vRealize Automation 8.x with vRealize Operations Manager , you can retrieve cloud zones into vRealize Operations Manager . The **Cloud Zones** option is hidden from the user until the integration with vRealize Automation 8.x is enabled from the integration page under **Administration** > **Management**.

The Cloud Zones option is enabled in vRealize Operations Manager , only if the following conditions are met.

■   vRealize Automation 8.x instance is integrated successfully in vRealize Operations Manager **Administration** > **Management**>**Integrations**.

■   vRealize Automation 8.x objects are discovered in vRealize Operations Manager .

■   vRealize Automation 8.x accounts and vRealize Operations vCenter Cloud Accounts are synchronized.

All the Cloud Zone objects which are existing in vRealize Automation 8.x environment, are discovered in vRealize Operations Manager . Cloud zones, whose dependent clusters are not discovered in vRealize Operations Manager , are not represented in Capacity Overview, Reclaim, and Workload Optimization pages.

Cloud Zones List

You can view the list of cloud zones that exist in your environment. In this view, you can click a cloud zone to display all the resources and objects that are associated with the cloud account. When you click the Cloud Zone, you are directed to the standard object summary page of the cloud account.

Where You Find Cloud Zones

Select **Environment** in the menu and click **Cloud Zones** tab.

Cloud Zone Tab Options

Option	Description
Name	Displays the name of the selected cloud zone.
Cloud Account	Displays the cloud accounts associated with the cloud zone.
Resources	Displays the cloud account resources associated with the cloud zone.
	Note If the resource field is empty, it means vRealize Operations Manager does not have a corresponding vCenter Cloud Account for that associated Cloud Zone. Add a new vCenter Cloud Account manually or use the Import Cloud Account option from the Cloud Account page.
Capability Tags	Displays the capability tags associated wit the cloud zone.

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.

- Support for vSAN datastores in workload optimization with cross-cluster rebalance actions.

  - You can move VMs from one vSAN datastore to another vSAN datastore.

- You can optimize the container if all the vSAN clusters are not in resync state.

- VMs with different storage policies for each disk or VMs with different types of storage for each disk will not be moved.

- You can generate a rebalance plan only if sufficient disk space is available at the destination vSAN datastore (The vSAN datastore slack space will also be considered).

- The storage policy assigned to the VM will be considered during the workload optimization (Compatibility check is performed against the storage policy).

- VM migration from vSAN datastore to vSAN stretched clusters is not supported.

Configure a vSAN Adapter Instance

When configuring an adapter instance for vSAN, you add credentials for a vCenter Server. In the earlier versions of vRealize Operations Manager , the vSAN solution was installed as part of the vRealize Operations Manager installation. Now, in case of a new installation the vSAN solution is pre-bundled as part of vRealize Operations Manager OVF, you must install the vSAN solution separately.

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.

You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to monitor the vSAN and Storage Devices.

To know how to install the Native Management Packs, see Solutions Repository.

Procedure

1   In the menu, select **Administration** and then select **Solutions > Cloud Accounts** from the left panel.

2   From the **Cloud Accounts** page, select the vCenter Server instance from the list and then click the **vSAN** tab.

3   To use the vCenter Server for enabling vSAN, move the **vSAN configuraton** option to the right.

Note   Once vSAN adapter instance is enabled and saved, the enable vSAN configuration option is not visible.

4   The credentials provided for the vCenter Server instance are also used for vSAN adapter instance. If you do not want to use these credentials, you can click **Use alternate credentials** option.

a   Click the plus sign next to the Credential field and enter the details in the **Manage Credentials** dialog box.

b   Enter the credential name, vCenter user name, and password and click **OK**.

5   Choose **Enable SMART data collection**, to enable SMART data collection for physical disk devices.

6   Click **Add**.

The vSAN configuration is enabled for the cloud account.

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

8   Accept the vCenter Server security certificate.

9   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. 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 > Adapter Instances > vSAN Adapter Instance > <*User_Created_Instance*>**.

Table 1-66. 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** .

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.

4   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.

5   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
     ]
   }
```

vSAN Log Analytics Enhancements

When vRealize Operations Manager is integrated with vRealize Log Insight, you can view and troubleshoot vRealize Log Insight object issues within vRealize Operations Manager . Earlier you could troubleshoot issues related only to vCenter objects, but now you can troubleshoot issues related to vSAN also.

The enhancements to vSAN log analytics include use of specific queries to retrieve log information for the following vSAN objects:

- vSAN Cluster

- Witness Host

- Disk Group

- Cache Disk

- Capacity Disk

Where You Find vSAN Object Logs

Navigate to the vSAN Object Details page, and click the **Logs** tab.

Note If you are not logged in to vRealize Log Insight, then vRealize Operations Manager prompts you to log in to vRealize Log Insight with your login credentials.

vRealize Operations Manager uses special queries for each object type. Using the special queries for vSAN objects, you can perform the following actions:

- View interactive analytics for the selected vSAN object.

- Retrieve log details for the vSAN object.

- Analyze and troubleshoot issues related to the vSAN object.

vRealize Network Insight

The vRealize Network Insight adapter enables integration of vRealize Operations Manager with vRealize Network Insight. VMware vRealize Network Insight provides network visibility and analytics to minimize risk during application migration, optimize network performance, manage and scale VMware NSX-T, VMware NSX for vSphere, vCenter on VMware Coud on AWS, VMware SD-WAN by VeloCloud, and Kubernetes deployments.

This adapter gets problem events from vRealize Network Insight and publishes the alerts in vRealize Operations Manager. Alerts are mapped correctly to the common objects between vRealize Network Insight and vRealize Operations Manager. Common objects supported in this adapter are vCenter Server, VMware NSX-T, and VMware NSX for vSphere. For the common objects, vRealize Operations supports launch-in-context to vRealize Network Insight. This allows the user to perform deep network troubleshooting with the vRealize Network Insight as the context.

The vRealize Network Insight adapter only supports vRealize Network Insight versions 5.2 and above. The vRealize Network Insight adapter can be installed and configured with On-prem version of vRealize Operations Manager or cloud version of vRealize Operations Cloud. The vRealize Network Insight adapter does not support cross platform configuration, it should be On-prem vRealize Operations Manager to On-prem vRealize Network Insight and vRealize Operations Cloud to vRealize Network Insight Cloud.

Configuring vRealize Network Insight

Configure an instance of the vRealize Network Insight in vRealize Operations Manager.

Prerequisites

As vCenter and NSX-T are native vRealize Operations Manager Management Packs, ensure that you have installed the latest NSX for vSphere Management Pack if you have NSX for vSphere data source configured in vRealize Network Insight.

Procedure

1 On the menu, click **Administration**.

2 In the left pane, expand **Management** and click **Integrations**.

3 Under Integrations, click the verticle ellipse next to VMware vRealize Network Insight and click **Configure**.

4 Configure the adapter instance.

Option	Description
VRNI FQDN/IP	The FQDN or the IP address of vRealize Network Insight.
Credential	Select and add the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the environment of this management pack, click the plus sign.
	■ **Credential Kind**. Select and configure the Credential Type. You can select either the Local, LDAP, or vIDM network insight credentials.
	Note This Management pack supports only the Local, LDAP, and vIDM users that are added in the **User Management** settings of vRealize Network Insight.
	■ Local - Network Insight Credentials. Enter the credential name, user name of the local user configured in vRealize Network Insight, and password for that user.
	■ LDAP - Network Insight Credentials. Enter the credential name, LDAP domain configured in vRealize Network Insight, LDAP user name, and LDAP password for that LDAP user.
	■ vIDM - Network Insight Credentials. Enter the credential name, vIDM FQDN/IP integrated with vRealize Network Insight, vIDM user name, and vIDM password for that vIDM user.
	Credentail Name. Credential Name.
Collector / Group	Select the required collector group.
Validate Connection	Test Connection should be successful.

5 The vRealize Network Insight instance collects events based on common data sources between vRealize Operations Manager and vRealize Network Insight. When you disable the **Import problem events as based on common data sources** option, all the events are imported into the vRealize Operations Manager.

6 You can collect user-defined events of vRealize Network Insight as notifications in vRealize Operations Manager. To do so, enable the **Import User defined events as Notifications**.

7 Select the severity of the problem events you want to import. By default, all the problem events with moderate and critical severities are imported.

8   Click **Add**.

The vRealize Network Insight instance is added to the list.

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.

Note   You cannot run the vRealize Application Remote Collector agent on the same VM as the End Point Operations Management agent.

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 1-67. 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
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	`epops-agent-x86-64-linux-version.rpm`
32bit Operating System	`epops-agent-x86-linux-version.rpm`
No Arch	`epops-agent-noarch-linux-version.rpm`

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	`epops-agent-x86-64-linux-version.tar.gz`
32bit Operating System	`epops-agent-x86-linux-version.tar.gz`
No Arch	`epops-agent-noJRE-version.tar.gz`

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 1-68. 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 1-69. 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.

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 ManagerServer 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.

Property	Property Definition
`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 `username` 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 thevRealize 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`	Provides details about the server certificate to trust.
	This parameter 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.

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

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

Operating System	Command
Linux	`./bin/ep-agent.sh set-property PropertyKey PropertyValue`
Windows	`./bin/ep-agent.bat set-property PropertyKey PropertyValue`

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	`ep-agent.sh start`
Windows	Install the Windows service for the agent, then run the `it: ep-agent.bat install ep-agent.bat start` command.
	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 anEnd 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.

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 `localhost`. 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 `thumbprint` 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 `agentManager` permissions.
Enter your server password	Enter the password for the specified vRealize Operations Manager. Do not store the password in the `agent.properties` 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+Dhp7213XQiyvtwI4tMlbGJfZMBPG23KnsUWu3OKrW35gB+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 our `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 our `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

Attribute	Description
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_Admistrator@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.procnet 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. On a freshly deployed instance of vRealize Operations Manager, before you register the End Point Operations Management agent, you must also manually activate the management pack from **Administration > Solutions > Repository > Operating Systems/Remote Service Monitoring**.

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	`ep-agent.sh setup`
Windows	`ep-agent.bat setup`

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.

Option	Description
dump	Runs a thread dump for the agent process, and writes the result to the `agent.log` file in `AgentHome/log`.
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 `start` after running `install`.
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 `agent.log` file in `AgentHome/log`.
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 thevCenter 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	1 Run `epops-agent remove`. 2 Remove the agent token and the `data` folder. 3 Run `epops-agent install`. 4 Run `epops-agent start`.

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 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 in 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 zxfp $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 zxfp $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	Supply the PTQL query in the form: `Class.Attribute.operator=value`.
	
	For example, `Pid.PidFile.eq=/var/run/sshd.pid`.
	
	Where:
	
	■ `Class` is the name of the Sigar class without the Proc prefix.
	
	■ `Attribute` is an attribute of the given Class, index into an array or key in a Map class.
	
	■ `operator` is one of the following (for String values):
	
	   ■ `eq` Equal to value
	
	   ■ `ne` Not Equal to value
	
	   ■ `ew` Ends with value
	
	   ■ `sw` Starts with value
	
	   ■ `ct` Contains value (substring)
	
	   ■ `re` Regular expression value matches
	
	Delimit queries with a comma.

Windows Service	Monitor an application that runs as a service under Windows. To configure it, you supply its Service Name in Windows. To determine the Service Name: 1 Select **Run** from the Windows Start menu. 2 Type `services.msc` 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	Configure vRealize Operations Manager to periodically run a script that collects a system or application metric.

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 Optionsand 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 `netservices` plug-in descriptor default values are:

- `port`: 80

- `sslport`: 443

HTTP Configuration Options

Table 1-70. ssl Option

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

Table 1-71. 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 1-72. 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 1-73. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
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 1-74. 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 1-75. 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 1-76. hostheader Option

Option Information	Value
Description	Host Header
Default	none
Optional	true
Type	N/A
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 1-77. 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 1-78. pattern Option

Option Information	Value
Description	Response Match (substring or regex)
Default	none
Optional	true
Type	N/A

Table 1-78. pattern Option (continued)

Option Information	Value
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 1-79. proxy Option

Option Information	Value
Description	Proxy Connection
Default	none
Optional	true
Type	N/A
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 1-80. 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 1-81. Credential Option

Option Information	Value
Description	Username
Default	N/A
Optional	true
Type	N/A

Table 1-81. Credential Option (continued)

Option Information	Value
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 1-82. 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

Table 1-83. 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 1-84. 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.

Table 1-84. port Option (continued)

Option Information	Value
Optional	false
Type	N/A
Notes	The port on which the service listens.
Parent Schema	sockaddr

Table 1-85. 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 1-86. 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
```

NSX-T

The NSX-T adapter allows you to retrieve alerts and findings from NSX-T to vRealize Operations Manager.

The NSX-T adapter supports adapter configuration using vIDM for NSX-T versions 3.0 and above. The roles and permissions associated with the vIDM users collecting the NSX-T adapter data is:

Roles	Permissions
Enterprise Admin	Collect all data.
VPN admin	Collect only Management appliance and NSX cluster data.
Network engineer	■ Collect all the NSX-T resources except the Load Balancer and collect limited routers data. Router data collected: ■ Tier 0 router connected to logical switch. ■ Tier 1 router created from vCloud Director.
■ Security Engineer ■ Security Operator ■ Auditor	Collect all data except the load balancer.
■ LB Admin ■ LB Auditor ■ Netxpartner Admin	Cannot collect any data.

Configuring the NSX-T Adapter

Procedure

1 On the menu, click **Administration** and in the left pane, click **Solutions > Other Accounts**.

2 On the Other Accounts page, click **Add Account**.

3 On the Account types page, click **NSX-T Adapter**.

4 Enter a display name and description for the NSX-T account.

   ▪ Name. Enter the name for the NSX-T instance as you want it to appear in vRealize Operations Manager.

   ▪ Description. Enter any additional information that helps you manage your instances.

5 Virtual IP/NSX-T Manager. Enter the FQDN, the IP address, or the Virtual IP of the NSX-T manager.

6 Select the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the NSX-T environment , click the plus sign.

   ▪ Credential Name. The name by which you are identifying the configured credentials.

   ▪ User Name. The user name of the NSX-T instance.

   ▪ Password. The password of the NSX-T instance.

7 Determine which vRealize Operations Manager collector or collector group is used to manage the account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

8 Click **Validate Connection** to validate the connection.

9 To save the configurations, click **Save This SDDC**.

10 Click **Add**.

   The adapter instance is added to the list.

What to do next

Verify that the adapter is configured and collecting data.

Configuring Alerts and Actions

<div style="text-align: right; font-size: 3em; color: #999;">2</div>

In vRealize Operations Manager , alerts and actions play key roles in monitoring the objects.

This chapter includes the following topics:

- Triggered Alerts
- Types of Alerts
- Alert Information
- Configuring Alerts
- Configuring Actions

Triggered Alerts

The **Triggered Alerts** page is a list of all the alerts generated in vRealize Operations Manager . Use the alert list to determine the state of your environment and to begin resolving problems.

How the Triggered Alerts Page Works

By default, only active alerts are initially listed, and the alerts are grouped by Time. Review and manage the alerts in the list using the toolbar options. Select multiple rows in the list using Shift+click, Control+click.

To see the alert details, click the alert name. The alert details appear on the right, including the symptoms triggered by the alert. The system offers recommendations for addressing the alert and link to run the recommendation. A Run Action button may appear in the details. Hover over the button to learn what recommendation is performed if you click the button. Alternatively, you can view the **Run** button and the **Suggested Fix** in the Alerts data grid. You can filter by alerts that have the Run option enabled and perform the recommended task to address the alert from the Alerts data grid. Click the small box on the lower left of the alert list to include the **Suggested Fix** and **Run** columns in the data grid.

Click the name of the object on which the alert was generated to see the object details, and access additional information relating to metrics and events.

If you migrated alerts from a previous version of vRealize Operations Manager , the alerts are listed with a cancelled status and alert details are not available.

Where You Find the All Alerts Page

In the menu, click **Alerts**.

Triggered Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to sort the alert list and to cancel, suspend, or manage ownership. Use the data grid to view the alerts and alert details.

Select an alert from the list to enable the Actions menu:

Table 2-1. Actions Menu

Option	Description
Cancel Alert	Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list. Cancel alerts when you do not need to address them. Canceling an alert does not cancel the underlying condition that generated it. Canceling alerts is effective if the alert is triggered by fault and event symptoms, because these symptoms are triggered again only if subsequent faults or events occur on the monitored objects. If the alert was generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again.
Delete Canceled Alerts	Delete cancelled (inactive) alerts by doing a group selection or by individually selecting alerts. The option is disabled for active alerts.
Suspend	Suspend an alert for a specified number of minutes. You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object. The user who suspends the alert becomes the assigned owner.
Assign to	Assign the alert to a user. You can search for a specific username and click **Save** to assign the alert to the slected user.
Take Ownership	As the current user, you make yourself the owner of the alert. You can only take ownership of an alert, you cannot assign ownership.
Release Ownership	Alert is released from all ownership.
Go to Alert Definition	Switches to the Alert Definitions page, with the definition for the previously selected alert displayed.

Table 2-1. Actions Menu (continued)

Option	Description
Disable...	Offers two options for disabling the alert:
	Disable the alert in all policies: this disables the alert for all objects for all the policies.
	Disable Alert in Selected Policies: this disables the alert for objects having the selected policy. Note that this method works only for objects with alerts.
Open an external application	Actions you can run on the selected object.
	For example, Open Virtual Machine in vSphere Client.

Table 2-2. Group By Options

Option	Description
None	Alerts are not sorted into specific groupings.
Time	Group alerts by time triggered. This is the default option. You can also group by 1 hour, 4 hours, Today and Yesterday, days of current week, Last week and Older.
Criticality	Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the "All Alerts Data Grid Options" table, below.
Definition	Group alerts by definition, that is, group like alerts together.
Object Type	Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together.

Table 2-3. All Filters

All Filters	Descriptions
Filtering options	Limit the list of alerts to those matching the filters you choose.
	For example, you might have chosen the Time option in the Group By menu. Now you can choose Status -> Active in the all Filters menu, and the All Alerts page displays only the active alerts, ordered by the time they were triggered.
Selected Options (see also the Group By and All Alerts Data Grid tables for more filter definitions:)	
Owner	Name of operator who owns the alert.
Impact	Alert badge affected by the alert. The affected badge, health, risk, or efficiency, indicates the level of urgency for the identified problem.

Table 2-3. All Filters (continued)

All Filters	Descriptions
Control State	State of user interaction with the alert. Possible values include: ■ Open. The alert is available for action and has not been assigned to a user. ■ Assigned. The alert is assigned to the user who is logged in when that user clicks **Take Ownership**. ■ Suspended. The alert was suspended for a specified amount of time. The alert is temporarily excluded from affecting the health, risk, and efficiency of the object. This state is useful when a system administrator is working on a problem and does not want the alert to affect the health status of the object.
Object Type	Type of object on which the alert was generated.
Updated On	Date and time when the alert was last modified. An alert is updated whenever one of the following changes occurs: ■ Another symptom in the alert definition is triggered. ■ Triggering symptom that contributed to the alert is canceled.
Canceled On	Date and time when the alert canceled for one of the following reasons: ■ Symptoms that triggered the alert are no longer active. Alert is canceled by the system. ■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions are disabled in the policy that is applied to the object. ■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions were deleted. ■ Alert definition for this alert is disabled in the policy that is applied to the object. ■ Alert definition is deleted. ■ User canceled the alert.
Action	Choose **Yes** to filter based on alerts that have the **Run** option enabled. Choose **No** to filter based on alerts that have the **Run** option disabled.

The Alerts data grid provides the list of generated alerts used to resolve problems in your environment. An arrow in each column heading orders the list in ascending or descending order.

Table 2-4. Triggered Alerts Data Grid

Option	Description
Criticality	Criticality is the level of importance of the alert in your environment.
	The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.
	The possible values include:
	■ Critical
	■ Immediate
	■ Warning
	■ Information
Alert	Name of the alert definition that generated the alert.
	Click the alert name to display the alert details to the right.
Triggered On	Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name.
	Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object.
Created On	Date and time when the alert was generated.
Status	Current state of the alert.
	Possible values include Active or Canceled.
Alert Type	Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/ Hypervisor, Hardware, Storage, and Network.
Alert Subtype	Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration.
Importance	Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm.
Suggested Fix	Displays the recommendation to address the alert.
Action	Click this button to perform the recommendation to address the alert.

Types of Alerts

Alerts in vRealize Operations Manager are of three types. The alert type determines the severity of the problem.

Health Alerts

The health alert list is all the generated alerts that are configured to affect the health of your environment and require immediate attention. You use the health alert list to evaluate, prioritize, and immediately begin resolving the problems.

Risk Alerts

The risk alerts list is all the generated alerts that are configured to indicate risk in your environment. Address risk alerts in the near future, before the triggering symptoms that generated the alert negatively affect the health of your environment.

Efficiency Alerts

The efficiency alerts list is all the generated alerts that are configured to indicate problems with the efficient use of your monitored objects in your environment. Address efficiency alerts to reclaim wasted space or to improve the performance of objects in your environment.

Alert Information

When you click an alert from the all alerts list, the alert information appears on the right. View the alert information to see the symptoms which triggered the alert, recommendations to fix the underlying issue, and troubleshoot the cause of the alert.

How You View the Alert Information

- In the menu, click **Alerts**. Click an alert from the alert list.

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. Click the object and then the **Alerts** tab.

- In the menu, select Search and locate the object of interest. Click the object and then the **Alerts** tab.

The alert description is hidden when you open the alert information. Click **View Description** to see the description of the alert. View the time stamp of when the alert started, and when it was updated, below the alert title.

Alert Details Tab

Section	Description
Recommendations	View recommendations for the alert. Click **<** or **>** to cycle through the recommendations. To resolve the alert, click the **Run Action** button if it appears.
Other Recommendations	Collapse the section to view additional recommendations. See the links in the **Need More Information?** section to view additional metrics, events, or other details that appear as a link.
Symptoms	View the symptoms that triggered the alert. Collapse each symptom to view additional information.
Notes	Enter your notes about the alert and click **Submit** to save.
Close	Click the X icon to close the alert details tab.

Related Alerts Tab

The **Related Scope** displayed on the right, shows the objects that are one level above and one level below the object on which the alert was triggered. This topology is fixed. You cannot change the scope in the **Related Alerts** tab.

On the right, you can see the following:

- If the same alert was triggered on the object in the past 30 days. This helps you understand if this is a recurring problem or something new.

- If the same alert was triggered on other peers in the same environment, in the past 30 days. This helps you do a quick peer analysis to understand if others are impacted with the same problem.

- All the alerts triggered in the current topology. This helps you investigate if there are other alerts upstream or downstream in the environment which are impacting the health of the object.

Potential Evidence Tab

See the **Potential Evidence** tab for potential evidences around the problem, and to arrive at the root cause. This tab displays events, property changes, and anomalous metrics potentially relevant to the alert. The time range and the scope are fixed. To modify the scope or the time range and investigate further, click **Launch Workbench**. This runs the troubleshooting workbench.

The time range that is displayed in the potential evidence tab is two hours and thirty minutes before the alert was triggered. vRealize Operations Manager looks for potential evidences in this time range.

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.

Metric / Super Metric Symptom Definitions

The Metric / Super Metric Symptom Definitions is a list of the metric-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined metric threshold triggering states and determine if you want to add, edit, or clone symptoms.

Where You Find Metric / Super Metric Symptoms

To manage symptoms based on metrics and super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-5. Metric / Super Metric Symptoms Options

Option	Description
Toolbar options	Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.
	■ Add. Add a symptom definition.
	Click the horizontal ellipsis to perform the following actions.
	■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge.
	■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge.
	■ Clone. Create a copy of the selected symptom definition.
	■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
All Filters	Limits the list to symptoms matching the filter. You can also sort on the columns in the data grid.
Quick Filter (Name)	Limits the list based on the text you type.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.

Table 2-5. Metric / Super Metric Symptoms Options (continued)

Option	Description
Metric Key	Text string that is used as a reference key for the metric. You can use the metric key to locate additional information about how the system statistics are derived from the metric.
Operator	Operator used to compare the current value to the threshold value, and trigger the symptom.
Threshold	Triggering threshold for the symptom. The threshold and the operator combine to set the point at which the symptom is triggered.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Metric and Supermetric Symptoms Definition Workspace

You define metric and super metric symptoms, which are based on collected operational or performance values, so that you can create one or more of the symptoms that you can add to an alert definition in vRealize Operations Manager . When a symptom is triggered, you use the symptoms to evaluate alerts or troubleshoot other problems.

How Metric Symptom Definitions Work

A metric or super metric symptom is triggered when a metric is compared to the configured static or dynamic thresholds, and the symptom condition is evaluated as true. If the symptom is based on a static threshold, the metric is compared based on the configured operator and the provided numeric value. If the symptom is based on a dynamic threshold, the metric is compared based on whether the current value is above, below, or abnormal compared to the calculated trend value.

Where You Find the Metric Symptom Definition Workspace

To define symptoms based on metrics or super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric / Property**. Click **Add** to define a metric-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-6. Symptoms Workspace Options for Metrics and Super Metrics

Option	Description
Metric Explorer	Components that you use to locate your metrics or super metrics for which you are creating symptoms.
Base Object Type	Object against which the symptom is evaluated. Based on the select object type, the list of available metrics displays only the metrics applicable to the object type.
Select Resource	If a metric or supermetric is not listed in the common metric or supermetric list, based on the selected based object type, use Select Resource to inspect the metrics or supermetrics of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a metric or supermetric for a specific object, the symptom definition is applicable to all objects with that metric or supermetric in your environment.
Search	Use a word search to limit the number of items that appear in the list.
Metric list	List of metrics for the selected base object type.
Symptom definition workspace	Click and drag the metric to the right pane. You can define symptoms based on static or dynamic thresholds.
Threshold	Determines if the symptom is static or dynamic. ■ Static thresholds are fixed values that trigger symptoms as true. You can configure one threshold for each symptom. You can also create multiple symptoms for multiple thresholds. For example, configure one symptom where the CPU use is greater than 90 percent and another where the CPU usage is less than 40 percent. Each is a separate symptom and can be added individually to an alert definition. ■ Dynamic thresholds are based on vRealize Operations Manager trended data where the triggering value is determined through the analytics. If the current value of the metric or super metric does not fall in the trended range, the symptom is triggered.

Table 2-6. Symptoms Workspace Options for Metrics and Super Metrics (continued)

Option	Description
Static Threshold configuration options	If you select Static Threshold, configure the options for this threshold type.
	■ Operator. Determines how the value you specify in the value text box is compared to the current value of the metric or super metric when the symptom is evaluated.
	■ Value. Value that is the triggering threshold.
	■ Criticality level. Severity of the symptom when it is triggered.
	■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.
	■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.
	■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.
	■ Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered.
	■ Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking **Select Object** next to the **Metrics** text box.
Dynamic Threshold configuration options	If you select Dynamic Threshold, configure the options for this threshold type.
	■ Threshold trend. Relationship of the current value to trended range based on the following options:
	■ Above. If current value is above trended range, the symptom is triggered.
	■ Below. If the current value is below the trended range, the symptom is triggered.

Table 2-6. Symptoms Workspace Options for Metrics and Super Metrics (continued)

Option	Description
	- Abnormal. If the current value is either above or below the trended range, the symptom is triggered. - Criticality level. Severity of the symptom when it is triggered. - Symtom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and viewing triggered symptoms. - Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered. - Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking **Select Object** next to the **Metrics** field.

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.

Property Symptoms Definitions

The Property Symptom Definitions is a list of the property-based symptoms in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined property triggering states and determine whether to add, edit, or clone symptoms.

Where You Find Property Symptoms

To manage symptoms based on properties, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-7. Property Symptoms Definitions Options

Option	Description
Toolbar options	Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click. ■ Add. Add a symptom definition. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
All Filters	Limits the list to symptoms matching the filter. You can also sort on the columns in the data grid.
Quick Filter (Name)	Limits the list based on the text you type.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Property	Text string that is used as a reference key for the property. You can use the property to locate additional information about the property.
Operator	Operator used to compare the threshold value to the current value.
Value	Text string that is the compared value for the property.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Property Symptoms Definition Workspace

You define property symptoms, which are based on collected configuration properties, so that you can add one or more symptoms to an alert definition in vRealize Operations Manager . You use the triggered symptoms to resolve alerts or troubleshoot other problems.

How Property Symptom Definitions Work

A property symptom is triggered when the defined threshold is compared with the current property value and the comparison is evaluated as true.

Where You Find the Property Symptom Definition Workspace

To define symptoms based on metrics or super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-8. Symptoms Workspace Options for Properties

Option	Description
Property Selector	Components that you use to locate the properties for which you are creating symptoms.
Based Object Type	Object against which the symptom is evaluated. Based on the selected object type, the list of available properties displays only the properties applicable to the object type.
Select Resource	If a property is not listed in the common properties list, based on the selected based object type, use Select Resource to inspect the properties of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a property for a specific object, the symptom definition is applicable to all objects with that property in your environment.
Search	Use a word search to limit the number of items that appear in the list.
Property list	List of properties for the selected base object type.

Table 2-8. Symptoms Workspace Options for Properties (continued)

Option	Description
Symptom definition workspace	Drag the property to the right pane.
Property	The properties are configured values that are compared to the value you specify. You can configure a single property symptom or add multiple symptoms.
	For example, if you need an alert when a particular property, such as Memory Hot Add, is no longer at the value required, you can configure a symptom and add it to an alert definition.
	Configure the options:
	■ Operator. Determines how the value you specify in the value text box is compared to the current value of the property for an object when the symptom definition is evaluated.
	■ Value. Value that the operator evaluates.
	■ Criticality level. Severity of the symptom when it is triggered.
	■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.
	■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.
	■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

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.

Message Event Symptom Definitions

The Message Event Symptom Definitions is a list of the message event-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined message events and to determine if you want to add, edit, or clone symptoms.

Where You Find Message Event Symptoms

To manage symptoms based on message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Message Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-9. Message Event Symptoms Options

Option	Description
Toolbar options	Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click. ■ Add. Add a symptom definition. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to symptoms matching the filter.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Event Type	Defined event classification type.
Operator	Operator used to compare the message from the incoming event against the event message specified in the symptom.
Event Message	Text string that is compared to the message in the incoming event using the specified operator.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Message Event Symptoms Definition Workspace

Message event symptoms are based on message events received from a component of vRealize Operations Manager or from an external monitored system through the system's REST API. You

define message event systems so that you can create one or more of the symptoms that you can add to an alert definition.

How Message Event Symptom Definitions Work

A message event symptom is triggered when a message in an incoming event matches the text string in the symptom based on the specified operator.

Where You Find the Message Event Symptom Definition Workspace

To define symptoms based on message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Message Event** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-10. Symptoms Workspace Options for Message Events

Option	Description
Message Event Selector	Components that you use to create symptoms.
Based Object Type	Object against which the symptom is evaluated.
Select the Type of Event	Select the type of incoming event against which you are matching the events as they arrive. The incoming event must contain the following type and subtype combinations. - System Performance Degradation. - Change. - Environment Down. - Notifications.

Table 2-10. Symptoms Workspace Options for Message Events (continued)

Option	Description
Symptom definition workspace	Drag the event type to the right pane.
Message Event	The Message Event text string is compared to the message in the incoming event by using the specified operator. You can configure a single message event symptom or add multiple symptoms.
	For example, the VMware adapter sends a change event when the CPU limit for a virtual machine was changed from unlimited to 2 GHz. You can create a symptom to detect CPU contention issues as a result of this configuration change.
	Configure the options:
	■ Operator. Determines how the string that you specify in the event message text box is evaluated against the message in the event when the symptom definition is evaluated.
	■ Event message. String that the operator evaluates.
	■ Criticality level. Severity of the symptom when it is triggered.
	■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.
	■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.
	■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

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.

Fault Symptom Definitions

The Fault Symptom Definitions is a list of the fault-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined fault message events and to determine whether to add, edit, or clone symptoms.

Where You Find Fault Symptoms

To manage symptoms based on fault message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Fault** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-11. Fault Symptoms Definitions Options

Option	Description
Toolbar options	Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click. ■ Add. Add a symptom definition. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to symptoms matching the filter.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Fault	Selected fault based on object type.

Table 2-11. Fault Symptoms Definitions Options (continued)

Option	Description
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Fault Symptoms Definition Workspace

You define fault symptoms, which are based on events published by the monitored systems, so that you can add one or more symptoms to an alert definition. You use the triggered symptoms to resolve alerts or troubleshoot other problems in vRealize Operations Manager .

How Fault Symptom Definitions Work

A fault symptom is triggered when a fault is active on the base object because of the occurrence of any of the fault events selected in the symptom definition.

Where You Find the Fault Symptom Definition Workspace

To define symptoms based on fault message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Fault** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-12. Symptoms Workspace Options for Faults

Option	Description
Fault Selector	Components that you use to create symptoms.
Based Object Type	Object against which the symptom is evaluated.
Fault definitions	Select the fault definition for the selected base object type. Some object types do not have fault definitions, and other types have multiple definitions.

Table 2-12. Symptoms Workspace Options for Faults (continued)

Option	Description
Symptom definition workspace	Drag the fault definition to the right pane.
Fault symptom definition	The fault events are published events from monitored systems. You can configure a single fault event symptom or add multiple symptoms.
	For example, if your base object is host and you drag the Hardware sensor fault for unknown type fault definition, you then select one of two text strings indicating a fault. Configure the options:
	■ Fault event. Select one or more fault events that activate the fault. If you do not select a string, then any of the provided strings are evaluated.
	■ Criticality level. Severity of the symptom when it is triggered.
	■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.
	■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.
	■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

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.

Metric Event Symptom Definitions

The Metric Event Symptom Definitions is a list of the metric event-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined threshold triggering states for the metric events and to determine if you want to add, edit, or clone symptoms.

Where You Find Metric Event Symptoms

To manage symptoms based on metric events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Click the **Metric Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-13. Metric Event Symptom Definitions Options

Option	Description
Toolbar options	Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click. ■ Add. Add a symptom definition. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to symptoms matching the filter.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.

Table 2-13. Metric Event Symptom Definitions Options (continued)

Option	Description
Event Metric	Selected event metric based on object type.
Event Type	Specifies whether the metric was above, below, equal to, or not equal to the threshold set by the monitoring system.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Metric Event Symptoms Definition Workspace

You define metric event symptoms, which are based on reported violations of metric thresholds from monitored systems, so that you can create one or more of the symptoms that you can add to an alert definition in vRealize Operations Manager .

How Metric Event Symptom Definitions Work

A metric event symptom is triggered when vRealize Operations Manager receives a metric event for the metric and event type defined in the symptom. The event type specifies whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system.

Where You Find the Metric Event Symptom Definition Workspace

To define symptoms based on metric events, in the left pane, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Metric Event** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 2-14. Symptoms Workspace Options for Metric Events

Option	Description
Metric Explorer	Components that you use to create symptoms.
Based Object Type	Object against which the symptom is evaluated. Based on the select object type, the list of available metrics displays only the metrics applicable to the object type.
Select Resource	If a property is not listed in the common properties list, based on the selected based object type, use Select Resource to inspect the properties of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a property for a specific object, the symptom definition is applicable to all objects with that property in your environment.
Search	Use a word search to limit the number of items that appear in the list.

Table 2-14. Symptoms Workspace Options for Metric Events (continued)

Option	Description
Metric Event list	List of the metric events for the selected base object type.
Symptom definition workspace	Click and drag the metric to the right pane.
Metric Event	You can configure a single threshold or add multiple thresholds.
	For example, configure a symptom where, when the virtual machine CPU usage is above the threshold defined in the monitored system, the metric event is above the threshold on the system.
	Configure the options:
	■ Event type. Select whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system.
	■ Criticality level. Severity of the symptom when it is triggered.
	■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.
	■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.
	■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

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 2-15. 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.

Recommendations

Recommendations are probable solutions for an alert generated in vRealize Operations Manager . You can create a library of recommendations that include instructions to your environment administrators or actions that they can run to resolve an alert.

Where You Find Recommendations

To define recommendations, in the menu, click **Alerts** and then in the left pane, click **Configuration > Recommendations**.

You can also define recommendations when you create an alert definition.

Table 2-16. Recommendations Overview Options

Option	Description
Toolbar options	Use the toolbar options to manage your recommendations. ■ Add. Add a recommendation. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected recommendation. ■ Delete. Remove the selected recommendation. ■ Clone. Create a copy of the selected recommendation so that you can create a new recommendation that uses the current one. ■ Export and Import. Export the file as XML from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to recommendations matching the filter.
Description	Recommendation text as it appears when the alert is generated and the recommendation is presented.
Action	If the recommendation includes running an action, the name of the actions.
Alert Definitions	Displays the number of alert definitions assigned for a particular recommendation. Click this link to view the alert definitions assigned for a particular recommendation and click **Remove from all** to remove the selected recommendation from all alert definitions.
Defined By	Indicates whether the recommendation was created by a user or provided with a solution adapter.

Table 2-16. Recommendations Overview Options (continued)

Option	Description
Last Modified	Displays the date on which the recommendation was last modified.
Modified By	Displays the name of the user who last modified the recommendation.

Recommendation Workspace

You create recommendations that are solutions to alerts generated in vRealize Operations Manager . The recommendations are intended to ensure that your network operations engineers and virtual infrastructure administrators can respond to alerts as quickly and accurately as possible.

How the Recommendations Workspace Works

A recommendation is instructions to your users or actions that your users can perform to resolve an alert. The instructions can be links to useful Web sites or local runbooks, instructions as text, or actions that you can initiate from vRealize Operations Manager .

Where You Find Recommendations Workspace

To define recommendations,click **Alerts** and then in the left pane, click **Configuration > Recommendations**. Click **Add** to create a recommendation.

You can also define recommendations when you define alerts.

Table 2-17. Define Recommendation Options

Option	Description
Create a hyperlink	Enter text in the text box, select the text, and click the button to make the text a hyperlink to a Web site or local wiki page. You cannot modify a hyperlink. To change the link, delete the hyperlinked word and create a new link.
Enter text	Enter the description of what must be done to resolve the triggered alert. The description can include steps a user must take to resolve the alert or it might be instructions to notify a virtual infrastructure administrator. This is a text field.
Adapter Type	Select an adapter type from the drop-down list to narrow down the list of actions displayed in the Actions field.
Action	You can add an action as a method to resolve a triggered symptom or a generated alert. Actions must already be configured in vRealize Operations Manager . You must provide text in the text box to describe the action before you can save the recommendation.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`

- `Set CPU Count for VM Power Off Allowed`

- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

Alert Definitions

Alert definitions are a combination of symptoms and recommendations that you combine to identify problem areas in your environment and generate alerts on which you can act for those areas. You use the Alert Definitions to manage your vRealize Operations Manager alert library, and to add or modify the definitions.

Where You Find Alert Definitions

To manage your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**.

Table 2-18. Alert Definition Options

Option	Description
Toolbar options	Use the toolbar options to manage your alert definitions. ■ Add. Add an alert definition. Click the horizontal ellipsis to perform the following actions. ■ Edit. Modify the selected definition. ■ Delete. Remove the selected definition. ■ Clone. Create a copy of the selected definition so that you can customize it for your needs. ■ Export and Import. Export the selected definition so that you can import it on another vRealize Operations Manager instance.
Filtering options	Limits the list of alerts to those matching the filter you create. You can also sort on the columns in the data grid.
Name	Name of the alert definition, which is also the name of the alert that appears when the symptoms are triggered.
Adapter Type	Adapter that manages the selected base object type.
Object Type	Base object type against which the alert is defined.

Table 2-18. Alert Definition Options (continued)

Option	Description
Alert Type	Metadata that is used to classify the alert when it is generated. You define the value on the Alert Impact page of the workspace.
Alert Subtype	Subcategory of the alert type and is the metadata that is used to classify the alert when it is generated. You define the value on the Alert Impact page of the workspace.
Criticality	Severity of the alert when it is generated. The criticality includes the following possible values: ■ Symptom. Alert is configured to display symptom based criticality. ■ Critical ■ Immediate ■ Warning ■ Info
Impact	Alert is configured to affect the Health, Risk, or Efficiency badge.
Defined by	Indicates who added the alert definition. The alert can be added by an adapter, a user, or the vRealize Operations Manager system.
Last Modified	Displays the date on which the alert was last modified.

Alert Definition Workspace

The alert definition process includes adding symptoms that trigger an alert and recommendations that help you resolve the alert. The alert definitions you create with this process are saved to your vRealize Operations Manager Alert Definition Overview list and actively evaluated in your environment based on your configured policies.

How the Alert Definition Workspace Works

You use the workspace to build alert definitions as you create the definition, the name, description, base object, and the alert impact. You can create or reuse existing symptoms and recommendations as part of the alert definition. If you create symptoms and recommendations, you add them to the definition, and they are added to the symptom and recommendations content libraries for future use. You also enable policies and select notifications for the alerts.

Where You Create an Alert Definition

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**. Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition.

Alert Definition Workspace Options

An alert definition is identified by a name and description. The definition comprises a target object type that is monitored for the alert, the badge that the alert affects, the set symptoms that trigger the alert, the recommendations that might resolve the alert, the policies that are enabled for an alert and the notification setting for which you want to receive the alert.

- Alert Definition Workspace Add Alert Details

  The name, description, base object type, and other details of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager .

- Alert Definition Workspace Add Symptom Definitions

  The add symptom definitions options are the mechanisms you use to add existing symptoms or to create new symptoms for the alert definition. If the symptom that you need for an alert definition does not exist, you can create it from this workspace.

- Alert Definition Workspace Add Recommendations

  Recommendations are instructions you provide to your user so that they can resolve generated alerts. The recommendations might include actions.

- Alert Definition Workspace Select Policies

  A policy is a set of rules that you define. It allows you to analyze and display information about the objects in your environment.

- Alert Definition Workspace Select Notifications

  Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager .

Alert Definition Workspace Add Alert Details

The name, description, base object type, and other details of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager .

Where You Define the Alert Details

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click**Configuration > Alert Definitions**. Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition. In the workspace, on the right, enter the details of alert definition.

Table 2-19. Alert Definition Details

Option	Description
Name	Name of the alert as it appears when the alert is generated.
Description	Description of the alert as it appears when the alert is generated. Provide a useful description for your users.

Table 2-19. Alert Definition Details (continued)

Option	Description
Base Object Type	The object type against which the alert definition is evaluated and the alert is generated. The drop-down menu includes all of the object types in your environment. You can define an alert definition based on one object type.
Impact	Under Advanced Settings, select the badge that is affected if the alert is generated. You can select a badge based on the urgency of the alert. ■ Health. Alert requires immediate attention. ■ Risk. Alert should be addressed soon after it is triggered, either in days or weeks. ■ Efficiency. Alert should be addressed in the long term to optimize your environment.
Criticality	Severity of the alert that is communicated as part of the alert notification. Select one of the following values. ■ Info. Informational purposes only. Does not affect badge color. ■ Warning. Lowest level. Displays yellow. ■ Immediate. Medium level. Displays orange. ■ Critical. Highest level. Displays red. ■ Symptom Based. In addition to alert criticality, each symptom includes a defined criticality. Criticality of the alert is determined by the most critical of all of the triggered symptoms. The color is dynamically determined accordingly. It you negate symptoms, the negative symptoms to not contribute to the criticality of a symptom-based alert.
Alert Type and Subtype	Select the type and subtype of alert. This value is metadata that is used to classify the alert when it is generated, and the information is carried to the alert, including the alert notification. You can use the type and subtype information to route the alert to the appropriate personnel and department in your organization.

Table 2-19. Alert Definition Details (continued)

Option	Description
Wait Cycle	The symptoms included in the alert definition remain triggered for this number of collection cycles before the alert is generated.
	The value must be 1 or greater.
	This setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition is added to the wait cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that after all of the symptoms are triggered at the desired symptom sensitivity level, the alert is immediately triggered.
Cancel Cycle	The symptoms are cancelled for this number of collection cycles after which the alert is cancelled.
	The value must be 1 or greater.
	This setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition is added to the cancel cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of the alert definition to 1. This configuration ensures that after all of the symptom conditions disappear after the desired symptom cancel cycle, the alert is immediately canceled.

Click **Next** to add symptom definitions.

Alert Definition Workspace Add Symptom Definitions

The add symptom definitions options are the mechanisms you use to add existing symptoms or to create new symptoms for the alert definition. If the symptom that you need for an alert definition does not exist, you can create it from this workspace.

How the Add Symptom Definitions Options Work

You can select and add symptoms defined for the base object type, and you can add symptoms for related object types. As you add one or more symptoms, you create a symptom expression. If this expression is evaluated as true, then the alert is generated.

Add Symptoms Definitions Options

To add symptom definitions, you can drag the selected symptom in to the left pane. Use the workspace on the left to specify whether all or any of the symptoms or symptom sets must be true to generate an alert.

Table 2-20. Add Symptoms Selection Options

Option	Description
Select Symptom	Select the type of symptom definition that you are adding for the current Defined On object type. ■ Metric / Property. Add symptoms that use metric and property symptoms. These metrics are based on the operational or performance values, and configuration properties that vRealize Operations Manager collects from target objects in your environment. ■ Message Event. Add symptoms that use message event symptoms. These 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. ■ Fault Event. Add symptoms that use fault symptoms. These symptoms are based on events that monitored systems publish. 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. ■ Metric Event. Add symptoms that use metric event symptoms. These 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 . These 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. ■ Smart Early Warning. Add a symptom that uses a defined condition that is triggered when the number of anomalies on an object is over the trending threshold. This symptom represents the overall anomalous behavior of the object. Anomalies are based on vRealize Operations Manager analysis of the number of applicable metrics that violate the dynamic threshold that determines the normal operating behavior of the object. This symptom is not configurable. You either use it or you do not use it.
Defined On	Object that the symptom evaluates. As you create alert definitions, you can select or define symptoms for the base object type and for related object types, based on the object relationship hierarchy. The following relationships are object types as they relate to the alert definition base object type. ■ Self. A base object type for the alert definition. For example, host system. ■ Descendant. An object type that is at any level below the base object type, either a direct or indirect child object. For example, a virtual machine is a descendant of a host system. ■ Ancestor. An object type that is one or more levels higher than the base object type, either a direct or indirect parent. For example, a datacenter and a vCenter Server are ancestors of a host system. ■ Parent. An object type that is in an immediately higher level in the hierarchy from the base object type. For example, a datacenter is a parent of a host system. ■ Child. An object type that is one level below the base object type. For example, a virtual machine is a child of a host system.
Filter by Object Type	Available only when you select a Defined On value other than Self. Limits the symptoms to those that are configured for the selected object type based on the selected Defined On relationship.

Table 2-20. Add Symptoms Selection Options (continued)

Option	Description
Create New Symptom	If symptoms that you need for your alert do not exist, you can create them. Opens the symptoms definition dialog box. Not available for Smart Early Warning symptoms, which are predefined in the system.
All Filters	Filter the list of symptom definitions. This selection is available when Defined On is set to **Self**, or when it is set to another relationship and you select an object from the Filter by Object Type drop-down menu. ■ Symptom. Type text to search on the name of the symptom definitions. For example, to display all symptom definitions that have efficiency in their name, type `Efficiency`. ■ Defined By. Type text to search for the name of the adapter that defines the symptom definitions. For example, to display all symptom definitions provided by the vCenter Adapter, type `vCenter`. To display only user-defined symptom definitions, type the search term `User`. To clear a filter, click the double arrow icon that appears next to the filter name.
Quick filter (Name)	Search the list based on the symptom name.
Symptoms list	List of existing symptoms for the selected object type. To configure a symptom, drag it into the left workspace. To combine symptoms that are based on multiple levels in the hierarchy, select the new Defined On level and Filter by Object Type before you select and drag the new symptom to the workspace.

Use the workspace to configure the interaction of the symptoms and symptom sets.

Table 2-21. Symptom Sets in the Alert Definition Workspace

Option	Description
Trigger alert when {operator} of the symptom sets are true	Select the operator for all of the added symptom sets. Available only when you add more than one symptom set. ■ All. All of the symptom sets must be true before the alert is generated. Operates as a Boolean AND. ■ Any. One or more of the symptom sets must be true before the alert is generated. Operates as a Boolean OR.
Symptoms	The symptom sets comprise an expression that is evaluated to determine if an alert should be triggered. To add one or more symptoms from the symptom list to an existing symptom set, drag the symptom from the list to the symptom set. To create a new symptom set for the alert definition, drag a symptom to the landing area outlined with a dotted line.

Symptom sets	Add one or more symptoms to the workspace, define the points at which the symptom sets are true, and specify whether all or any of the symptoms in the symptom set must be true to generate the alert.
	A symptom set can include one or more symptoms, and an alert definition can include one or more symptom sets.
	If you create a symptom set where the Defined On object is Self, you can set the operator for multiple symptoms in the symptom set.
	If you create a symptom set where the Defined On object is a relationship other than Self, you can set the operator and modify the triggering threshold. To configure the symptom set criteria, you set the options.
	■ Value operator. Specifies how the value you provide in the value text box is compared to a number of related objects to evaluate the symptom set as true.
	■ Value text box. Number of objects of the specified relationship, based on the value type, that are required to evaluate the symptom set as true.
	■ Value type. Possible types include the following items:
	  ■ Count. Exact number of related objects meet the symptom set criteria.
	  ■ Percent. Percentage of total related objects meet the symptom set criteria.
	  ■ Any. One or more of the related objects meet the symptom set criteria.
	  ■ All. All of the related objects meet the symptom set criteria.
	■ Symptom set operator. Operator applied between symptoms in the symptom set.
	  ■ All. All of the symptoms must be true before the alert is generated. Operates as a Boolean AND.
	  ■ Any. One or more of the symptoms must be true before the alert is generated. Operates as a Boolean OR.
	When you include a symptom in a symptom set, the condition must become true to trigger the symptom set. However, you might want to configure a symptom set where the absence of a symptom condition triggers a symptom. To use the absence of the symptom condition, click the vertical ellipsis on the left of the symptom name and select **Invert Symptom**.

Table 2-21. Symptom Sets in the Alert Definition Workspace (continued)

Option	Description
	Although you can configure symptom criticality, if you invert a symptom, it does not have an associated criticality that affects the criticality of generated alerts.

Click **Next** to add recommendations.

Alert Definition Workspace Add Recommendations

Recommendations are instructions you provide to your user so that they can resolve generated alerts. The recommendations might include actions.

How Add Recommendations Works

Recommendations are information provided to users to resolve a problem when an alert is generated. You use the recommendation options to add existing information or to create solutions to alerts. If the recommendation that you need for an alert definition does not exist, you can create it from this workspace.

Add Recommendations Options

To add recommendations, you can drag the selected recommendation in to the left pane. Use the workspace on the left to to change the priority order.

Table 2-22. Add Recommendations Options in the Alert Definition Workspace

Option	Description
Create New Recommendation	If recommendations that you need to resolve the symptoms in the problem do not exist, you can create them.
All Filters	Filter the list of recommendations. ■ Description. Type text to search on the name of the recommendation. For example, to display all recommendations that have memory in their name, type `Memory`. ■ Defined By. Type text to search for the name of the adapter that defines the recommendation. For example, to display all recommendations provided by the vCenter Adapter, type `vCenter`. To clear a filter, click the double arrow icon that appears next to the filter name.
Quick filter (Name)	Limits the list based on the text you enter.

Table 2-22. Add Recommendations Options in the Alert Definition Workspace (continued)

Option	Description
List of available recommendations.	List of existing recommendations that you can drag to the workspace.
	Recommendations are instructions and, where possible, actions that assist you with resolving alerts when they are triggered.
Recommendation workspace	Add one or more recommendations to the workspace.
	If you add more than one recommendation, you can drag the recommendations to change the priority order.

Click **Next** to enable policies.

Alert Definition Workspace Select Policies

A policy is a set of rules that you define. It allows you to analyze and display information about the objects in your environment.

How the Select Policies Option Works

Policies define the settings that vRealize Operations Manager applies to your objects when it collects data from your environment. You can select the policies that you want to apply for a particular alert.

Select Policies Option

You can view the policy tree in the left pane and you can either select the default policy or any other policy from the tree. You can also customize thresholds for a policy by clicking the policy and editing the trigger value in the right pane.

Note If you create an alert without enabling any policies, then the alert remains inactive.

Click **Next** to select notifications.

Alert Definition Workspace Select Notifications

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager .

How the Select Notifications Option Works

You can send alert notifications for an alert by assigning the alert to a notification rule that you have set up.

Select Notifications Option

You can view the notification setting on the left pane and select the notification setting for which you want to receive the alert.

Click **Create** to create the alert. The new alert appears in the list of alert definitions.

Create a Simple Alert Definition

While troubleshooting, you can now quickly create an alert for a particular object type or a metric in a quick and efficient way.

You can create a simple alert definition from the following locations.

- In the **Home** page, click **Troubleshoot > Workbench** and select the metric for which you want to create an alert. You can create an alert from the **Potential Evidence** or the **Metrics** tab.

- In the **Alerts** page, click **Triggered Alerts**. Select an alert and click the **Potential Evidence** tab.

Procedure

1   Click the drop-down menu available in the right side of the widget and select the **Create an Alert Definition** option.

2   In the Create Alert Definition page, enter the **Name** and **Description** of the alert.

3   Set thresholds, criticality, and the number of wait cycles. Click **Show Advanced Settings** to set Wait Cycle and Cancel Cycle.

Note   The Object Type or Metric/Property are pre-selected and cannot be edited.

4   Click **Create**.

The new alert is created and the policy the object belongs to and its children policies are enabled for 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 **Configuration > 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 affect 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 <=50% in one alert definition and >=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.

Outbound Settings

You use the Outbound Settings to manage your communication settings so that you can send information to users or applications outside of vRealize Operations Manager .

How Outbound Settings Work

You manage your outbound options from this page, including adding or editing outbound plug-ins, and turning the configured plug-ins on or off. When enabled, the plug-in sends a message to users as email notifications, or sends a message to other applications.

Where You Find Outbound Settings

To manage your outbound settings, select **Administration** in the left pane, and click **Management > Outbound Settings**.

Table 2-23. Outbound Settings Options

Option	Description
Toolbar options	Use the toolbar options to manage your Outbound Plug-Ins. ■ Add. Opens the Outbound Plug-In dialog box where you configure the connection options for the instance. Select an existing plugin and click the vertical ellipsis to perform the following actions. ■ Edit. Modify the Outbound Plug-In instance details. ■ Delete. Removes the selected plug-in instance. ■ Enable or Disable. Starts or stops the plug-in instance. Disabling an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.
Instance Name	Name that you assigned when you created the plug-in instance.
Plug-In Type	Type of configured plug-in for the plug-in instance. The types of plug-ins vary depending on the solutions you added to your environment. The most common plug-in types include standard email, SNMP trap, log file, and REST.
Status	Specifies whether the plug-in is currently running.

Outbound Plug-Ins

Outbound plug-in settings determine how the supported external notification systems connect to their target systems. You configure one or more instances of one or more plug-in types so that you can send data about generated notifications outside of vRealize Operations Manager .

How Outbound Plug-Ins Work

You configure each plug-in with the required information, including destination locations, hosts, ports, user names, passwords, instance name, or other information that is required to send notifications to those target systems. The target systems can include email recipients, log files, or other management products.

Some plug-ins are included with vRealize Operations Manager , and others might be added when you add a management pack as a solution.

Where You Configure Outbound Settings

To add or edit an outbound plug-in, select **Administration** in the top pane, and click **Management > Outbound Settings**. Click **Add** to add a plug-in instance or select a plug-in, click the vertical ellipsis and select **Edit** to edit the existing plug-in.

Outbound Plug-In Configuration Options

The configuration options vary depending on which plug-in you select from the **Plug-In Type** drop-down menu.

To add outbound notification plug-in, see Add Outbound Notification Plug-Ins in vRealize Operations Manager .

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 2-24. 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, select the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only have to provide the instance name.
Log File plug-in	Yes To filter the log file alerts, you can either configure the file named `TextFilter.xml` 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
Slack plug-in	Yes
Service-Now Notification 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

- Add a REST plug-in for vRealize Operations Manager Outbound Alerts

- 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 only SMB version 2.1.

- 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 Service-Now Notification Plug-In for Outbound Alerts

  You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with vRealize Operations Manager . Service Now creates an incident whenever an alert is triggered in vRealize Operations Manager .

- Notifications - Add a Slack Plug-In for Outbound Notifications

  You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run vRealize Operations Manager using alert links to troubleshoot further.

- Sample Email Alert

  Here is a sample email for a newly created alert.

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 **Add**.

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.
Receiver Email Address	Receiver's 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 menu, click **Administration** and then in the left pane, click **Management**.

2 Click **Outbound Settings** and click **Add**.

3 From the **Plug-In Type** drop-down menu, select **Rest Notification plug-in**.

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.

Option	Description
Content Type	Specify the format for the alert output. ■ 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 **Add**.

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 only SMB version 2.1.

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 **Add**.

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.

Option	Description
Password	The password for the domain user account.
Network share root	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.
	You must enter an IP address. For example, `\\IP_address\ShareRoot`. 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.
	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.

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. See Schedule Reports Overview.

Sample Log File Plug-In Output
Here is a sample log file plug-in output.

```
AlertId :: 9fb52c9c-40f2-46a7-a005-01bf24ab75e6

AlertStatus :: Active

AlertControlState :: Open

AlertGenerateTime :: Wed May 06 06:26:05 UTC 2020 (UTC = 1588746365585)

AlertUpdateTime :: Wed May 06 06:26:05 UTC 2020 (UTC = 1588746365585)

AlertMessage :: 9027

AlertSummaryLink :: https://10.27.82.96/ui/index.action#/object/all/1b852a3c-bbdf-41df-a64d-
b40af9673b89/alertsAndSymptoms/alerts/9fb52c9c-40f2-46a7-a005-01bf24ab75e6

AlertType :: Storage - Performance

AlertCriticality :: 4

AffectedResourceId :: 1b852a3c-bbdf-41df-a64d-b40af9673b89
```

```
AffectedResourceName :: JNJ_6nodes_Large_HA_4_10.27.83.44


AffectedResourceKind :: VirtualMachine


AffectedResourceParentsNames ::
 VM Entity Status:PoweredOn:all
 DistributedVirtualPortgroup:VM-Network-VLAN-820
 VM Entity Status:PoweredOn:vc_evn-hs1-vc.company.com
 VMFolder:Discovered virtual machine
 HostSystem:evn1-hs1-0808.company.com


AffectedResourceAdapterInstanceResourceName ::
 CompanyAdapter Instance:vc_evn-hs1-vc.company.com


AlertOwner ::


Anomalies ::
 VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 25
 VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 15
 VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 30


Health ::
4.0
Risk ::
2.0
Efficiency ::
1.0
KPIFiring ::


AlertTrigger ::
 Resource                                        Message Info
Alarm Reason     Probability                Prediction Time
 VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44     HT above 30.5647619047619 > 30     HT
above        Unable to retrive value      Unable to retrive value


AlertRootCause ::
null
AlertRootCauseDetails :: null


AlertName :: Virtual machine disk I/O write latency is high


AlertDescription ::
Virtual machine disk I/O write latency is high
```

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 **Add**.

3   From the **Plug-In Type** drop-down menu, select **SNMP Trap Plugin**.

    The dialog box expands to include your SNMP trap settings.

4   Enter 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 SNMPv1 and SNMPv2c protocol.
Username	User name to configure SNMP trap settings in your environment. If the user name 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, AES256.
Privacy Password	Privacy password.
Engine ID	Engine ID serves as an identifier for the agent. It is used with a hashing function to generate localized keys for authentication and encryption of SNMP v3 messages.
	It is mandatory to specify the Engine ID when configuring the SNMP Trap plugin. If you do not add the Engine ID and save the SNMP Trap plugin instance, the field is auto-generated the next time you edit the settings.

6   Click **Test** to validate the connection.

Note   The Community and Username options are mutually exclusive. Define either one of them to avoid an error. If you add a user name, you can optionally define the Authentication Protocol and Authentication Password followed by the Privacy Protocol and Privacy Password. The privacy protocol and its password cannot be defined independent of the authentication protocol and its password.

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 **Add**.

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.

Option	Description
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.
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::CAPA
CITY,APPLICATION::COMPLIANCE,VIRTUALIZATION::AVAILABILITY,VIRTUALIZATION::PERFO
RMANCE,VIRTUALIZATION::CAPACITY,VIRTUALIZATION::COMPLIANCE,HARDWARE::AVAILABILI
TY,HARDWARE::PERFORMANCE,HARDWARE::CAPACITY,HARDWARE::COMPLIANCE,STORAGE::AVAIL
ABILITY,STORAGE::PERFORMANCE,STORAGE::CAPACITY,STORAGE::COMPLIANCE,NETWORK::AVA
ILABILITY,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.

Add a Service-Now Notification Plug-In for Outbound Alerts
You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with vRealize Operations Manager . Service Now creates an incident whenever an alert is triggered in vRealize Operations Manager .

Using Service-Now Notification Plug-In you can send alert notifications to the Service Now ticketing system to create incidents. The incident includes information like the Caller, Category, Subcategory, Business Service, and other attributes related to alerts.

Prerequisites

Ensure that you have log in credentials for Service-Now.

Ensure that you are assigned with IT Infrastructure Library (ITIL) role in Service Now.

Procedure

1   In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.

2   From the toolbar, click **Add** and from the **Plug-In Type** drop-down menu, select **Service-Now Notification Plug-in**.

    The dialog box expands to include your plug-in instance settings.

3   Enter an **Instance Name**.

4   Enter the Service Now URL.

    `https://dev22418.service-now.com/`

5   Enter the user name and password for Service Now.

6   Enter a value for the Connection Count.

    The connection count represents the maximum number of open connections allowed per node in vRealize Operations Manager .

7   To verify the specified paths, credentials, and permissions, click **Test**.

8   Click **Save**.

Results

This instance of the Service-Now Notifications plug-in is configured and running.

What to do next

When the plug-in is added, Configuring Notifications for creating incidents in Service-Now ticketing system.

Notifications - Add a Slack Plug-In for Outbound Notifications
You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run vRealize Operations Manager using alert links to troubleshoot further.

Prerequisites

Procedure

1   In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.

2 From the toolbar, click **Add** and from the **Plugin Type** drop-down menu, select **Slack Plugin**.

The dialog box expands to include your plug-in instance settings.

3 Enter an **Instance Name**.

4 Enter a value for the **Connection Count**.

The connection count represents the maximum number of open connections allowed per node in vRealize Operations Manager .

5 To verify the specified paths, credentials, and permissions, click **Test**.

6 Click **Save**.

Results

This instance of the Slack plugin is configured and running.

What to do next

When the plugin is added, Configuring Notifications for different slack channels.

Sample Email Alert

Here is a sample email for a newly created alert.

```
Alert Definition Name: Node is experiencing swapping due to memory pressure
Alert Definition Description: Node is experiencing swapping due to memory pressure
Object Name : vRealize Operations Manager Node-vRealize Cluster Node
Object Type : vC-Ops-Node
Alert Impact: risk
Alert State : warning
Alert Type : Application
Alert Sub-Type : Performance
Object Health State: info
Object Risk State: warning
Object Efficiency State: info
Control State: Open
Symptoms:
SYMPTOM SET - self
```

Symptom Name	Object Name	Object ID	Metric	Message Info
Node swap usage at Warning level	vRealize Operations Manager Node-vRealize Cluster Node	50ec874a-2d7d-4e78-98b1-afb26fd67e58	Swap|Workload	59.183 > 30.0

```
Recommendations:
Notification Rule Name: rule1
Notification Rule Description:
Alert ID : badc2266-935d-4fb9-8594-e2e71e4866fc
VCOps Server - vRealizeClusterNode

Alert details(link)
```

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 **Configuration > Notifications**.

2   Click **Add** to add a notification rule.

3   In the **Name** text box, enter 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, enter 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, enter 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 Set the **Notification Status**, you can either enable or disable a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.

7 Configure the scope of filtering criteria.

a   From the **Scope** drop-down menu, select **Object**.

b   Click **Select an Object** and enter the name of the object.

In this example, type `mmbhost`.

c   Locate and select the object in the list, and click **Select**.

8 Configure the Notification Trigger.

a   From the **Notification Trigger** drop-down menu, select **Impact**.

b   From the adjacent drop-down menu, select **Health**.

9 In the Criticality area, click **Critical**.

10 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.

11 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 a 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 **Configuration > Notifications**.

2 Click **Add** to add a notification rule.

3 In the **Name** text box, enter a name similar to `Virtualization Alerts for Availability`.

4 In the Method area, select **REST Notification Plugin** from the drop-down menu, and select the configured instance of the email plug-in.

5 Set the **Notification Status**, you can either enable or disable a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.

6 Configure the Notification Trigger.

    a From the **Notification Trigger** drop-down menu, select **Alert Type**.

    b Click **Select an Alert Type/Subtype** and select any alert types or subtypes under **Virtualization/Hypervisor Alerts Availability**.

7 In the Criticality area, click **Warning**.

8 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.

9 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.

Notifications

You use the Notifications page to manage your individual alert notification rules. The rules determine which vRealize Operations Manager alerts are sent to the supported target systems.

How Notifications Work

You add, manage, and edit your notification rules from this page. To send notifications to a supported system, you must configure and enable the settings for outbound alerts. The supported outbound notification plug-ins include the Standard Email Plug-In, REST plug-in, SNMP Trap plug-in, and the Log File plug-in.

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances.

Where You Find Notifications

To manage your notifications, in the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**.

Table 2-25. Notifications Options

Option	Description
Toolbar options	Use the toolbar options to manage your notification rules. ■ Add. Opens the Add Rule dialog box where you configure the filtering options for the notification rule. Click the horizontal ellipsis to perform the following actions. ■ Delete. Removes the selected rule. ■ Disable or Enable. Disables or enables the selected rule(s). ■ Export or Import. Export the selected notifications to a ".xml" file so that you can import it on another vRealize Operations Manager instance.
Rule Name	Name you assigned when you created the notification rule. Click the vertical ellipsis to perform the following actions. ■ Edit. Allows you to edit the selected rule. ■ Delete. Removes the selected rule. ■ Disable or Enable. Disables or enables the selected rule.
Instance	Name of the configured outbound alert instance for the notification rule. Instances are configured as part of the outbound alerts and can indicate different email servers or sender addresses for alert notifications.
Enabled	Displays if the rule is enabled or not.
Email Address	If the rule is for standard email notifications, the alert recipient email addresses are listed.
Object Name	If the rule specifies a notification for a particular object, the object name is listed.
Children	If the rule specifies a notification for a particular object and selected child objects, the child object types are listed.
Last Modified	Displays the date on which the rule was last modified.
Modified By	Displays the name of the user who last modified the rule.

Notification Rule

Notification rules determine which alerts are sent to the target systems. You configure one or more notification rules to limit the data that vRealize Operations Manager sends to systems or recipients.

How Notification Rules Work

Notification rules are filters that limit the data sent to external systems by using outbound alert plug-ins that are supported, configured, and running. Rather than sending all alerts to all your email recipients, you can use notification rules to send specific alerts. For example, you can send health alerts for virtual machines to one or more of your network operations engineers. You can send critical alerts for selected hosts and clusters to the virtual infrastructure administrator for those objects.

Before you can create and manage notification rules, you must configure the outbound alert plug-in instances.

You can configure one filtering selection, or you can configure as many selections as you need so that vRealize Operations sends only the required data to the target external system.

Where You Find Notification Rules

To manage your notifications, in the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.

Table 2-26. Notification Rule Configuration Selections

Selections	Description
Name	Name of the rule that you use to manage the rule instance.
Method	Includes plug-in type and the plug-in instance. If you are configuring notifications for standard email, you can add recipients and associated information. ■ Type of plugin. Select one of the configured outbound alert plug-in types: Standard Email, REST, SNMP Trap, Log File, Service-Now, and Slack Plugin. ■ Instance. Select the configured instance for the type of plug-in.
Method -Standard Email Plugin	Includes plug-in type and the plug-in instance. If you are configuring notifications for standard email, you can add recipients and associated information. ■ Recipients. Enter the email addresses of the individuals to whom you are sending email messages that contain alert notifications. If you are sending to more than one recipient, use a semicolon (;) between addresses. ■ Notify again. Number of minutes between notifications messages for active alerts. Leave the text box empty to send only one message per alert. ■ Max Notifications. Number of times to send the notification for the active alert. Leave the text box empty to send only one message per alert. ■ Delay to notify. Number of minutes to delay before sending a notification when a new alert is generated. For example, if the delay is 10 minutes and a new alert is generated, the notification is not sent for 10 minutes. If the alert is canceled in those 10 minutes, the notification is not sent. The notification delay reduces the number of notifications for alerts that are canceled during that time. ■ Description. Enter the text to include in the email message. For example, Attention Host Management team.

Table 2-26. Notification Rule Configuration Selections (continued)

Selections	Description
Method - Service-Now Notification Plugin	If you are configuring notifications for a Service-Now notification plug-in, you can add instances and associated information. ■ Caller. Enter the name of the person who reported the incident or who is affected by the incident. ■ Category. Specify the category to which the incident belongs. ■ Sub Category. Specify the sub category to which the incident belongs. ■ Business Service. Specify the business service of the incident. ■ Contact Type. Enter the contact type. ■ State. Enter the incident state in digits. ■ Resolution Code. Enter the resolution code for the incident. ■ Resolution notes. Enter the resolution notes for the incident. ■ On hold reason. Enter the reason as to why the incident is on hold. ■ Impact. Set the incident impact in digits. Impact measures the business criticality of the affected service. ■ Urgency. Set urgency for the incident in digits. Urgency defines the number of days taken to resolve an incident. ■ Priority. Enter the priority for the incident. Priority defines the sequence in which the incident must be resolved. ■ Assignment Group. Enter the assignment group for the incident. ■ Assigned To. Enter the details of the person to whom the incident is assigned. ■ Severity. Set the severity for the incident in digits. ■ Upon Approval. Specify the next steps to be taken upon incident approval. ■ Problem. Enter the details of the related problem if it exists. ■ Cause by change. Enter the change request which triggered the incident. ■ Change Request. Enter the details for the related change list if it exists.
Method - Slack Plugin	If you are configuring notifications for a Slack plugin, add the Webhook URL of Slack. For example, the Webhook URL is in the format: `https://hooks.slack.com/services/T00000000/B00000000/XXXXXXXXXXXXXXXXXXXXXXXX`. Create and authorize an app within Slack to obtain the Webhook URL. For details on creating and authorizing an app within Slack, refer to the Slack Documentation. Once you have created the notification rule, the alerts are displayed within that particular Slack channel with a link to the alert. Click the link to view the details of the alert in the Object Summary page.
Notification Status	Either enable or disable a notification setting. Disabling a notification will stop the alert notification for that setting and enabling it will activate it again.
Filtering Criteria	**Note** The Filtering Criteria and Advanced Filter sections are same for all the plugins.
Scope	General object type for which you are filtering the alert notifications. After you select the type, you select the specific instance. For example, if you select Object, you then select the specific object by name and determine whether to include any child objects.
Notification Trigger	Alert type and subtypes, impact, or definition that triggers the alert. After you select the trigger type, you configure the specific selections associated with the trigger type. For example, if you select Alert Definition, you then select the alert definition that limits the data to alerts with this definition. You can select multiple alert definitions as conditions for a notification to trigger.

Table 2-26. Notification Rule Configuration Selections (continued)

Selections	Description
Criticality	Defined criticality of the alert that results in the data being sent to an external system. For example, if you select Critical, then the data that is sent to the external system must also be labeled as critical.
Advanced Filters	
Alert States	Managed state of the alert, either opened, assigned, or suspended.
Alert Status	Current state of the alert, either canceled, updated, or new.
Collector/Group	Configured collectors in your environment. For example, in an environment where you manage multiple vCenter Server instances, you can select a collector for one instance. If you want to distribute email alert notifications between various groups which use different remote collectors, select **Default collector group**. This option filters alerts by the target collector group.

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**Configuration > Alert Definitions**.

2 Click **Add** 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   From the **Base Object Type** 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.

5   Click **Advanced Settings** 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, select **Performance** under **Virtualization/Hypervisor**.

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 **Next** and configure the symptoms.

2   Begin configuring the symptom set related to virtual machines CPU usage.

    a   From the **Select Symptom** drop-down menu, select **Metric / Property**.

    b   From the **Defined On** drop-down menu, select **Child**.

    c   From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.

    d   Click **Create New** 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 left.

    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 left.

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 **Next**.

2 Configure the symptom related to host systems for the virtual machines.

    a   From the **Select Symptom** drop-down menu, select **Metric / Property**.

    b   From the **Defined On** drop-down menu, select **Self**.

    c   Click **Create New** to add 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 left.

    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 left.

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 **Next** and add the recommended actions and instructions.

2   Click **Create New Recommendation** and select an action recommendation to resolve the virtual machine alerts.

    a   In the **Description** 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 **Create**.

3   Click **Create New Recommendation** 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 **Create**.

5 Click **Create New Recommendation** 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 the hyperlink icon.

   d   Paste the URL in the **Create a hyperlink** text box and click **OK**.

   e   Click **Create**.

6 In the **Alert Recommendation 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.

7 Click **Next** to select policies and view notifications.

8 Click **Create**.

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 **Add** 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 you 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 and then, click **Add** .

3 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.
```

4 Select **Default Policy** from the **Start with** drop-down menu.

5   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.

6   On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.

7   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 **Configuration > Notifications**.

2   Click **Add** 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 **Recipient(s)** 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 **Select Alert Definitions**.

    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 **Triggered 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 2-27. vRealize Operations Manager Actions Affected Objects

Action	Modified Object	Object Levels
Rebalance Container	Virtual Machines	- Data Center - Custom Data Center
Delete Idle VM	Virtual Machines	- Clusters - Host Systems - Virtual Machines
Set DRS Automation	Cluster	Clusters
Move VM	Virtual Machine	Virtual Machines
Power Off VM	Virtual Machine	- 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.	- Clusters - Host Systems - Virtual Machines
Power On VM	Virtual Machine	- Clusters - Host Systems - Virtual Machines
Delete Powered Off VM	Virtual Machine	- Clusters - Host Systems - Virtual Machines
Set Memory for VM and Set Memory for VM Power Off Allowed	Virtual Machine	- Clusters - Host Systems - Virtual Machines
Set Memory Resources for VM	Virtual Machine	- Clusters - Host Systems - Virtual Machines
Set CPU Count for VM and Set CPU Count for VM Power Off Allowed	Virtual Machine	- Clusters - Host Systems - Virtual Machines
Set CPU Resources for VM	Virtual Machine	- Clusters - Host Systems - Virtual Machines

Table 2-27. vRealize Operations Manager Actions Affected Objects (continued)

Action	Modified Object	Object Levels
Set CPU Count and Memory for VM and Set CPU Count and Memory for VM Power Off Allowed	Virtual Machine	■ Clusters ■ Host Systems ■ Virtual Machines
Delete Unused Snapshots for VM	Snapshot	■ Clusters ■ Host Systems ■ Virtual Machines
Delete Unused Snapshots for Datastore	Snapshot	■ Clusters ■ Datastores ■ Host Systems
Execute Script	Virtual Machine	■ Virtual Machine
Get Top Processes	Virtual Machine	■ Virtual Machine
Apply Guest User Mapping	vCenter Server	■ vCenter Server **Note** This action is deprecated and will be removed in the next release.
Clear Guest User Mapping	vCenter Server	■ vCenter Server **Note** This action is deprecated and will be removed in the next release.
Export Guest User Mapping	vCenter Server	■ vCenter Server **Note** This action is deprecated and will be removed in the next release.
Configure Included Services	Service Discovery Adapter Instance	■ Service Discovery Adapter Instance **Note** This action is deprecated and will be removed in the next release.

Actions Overview List in vRealize Operations Manager

Actions are the method you use to configuration changes on managed objects that you initiate from vRealize Operations Manager . These actions are available to add to alert recommendations.

How the Actions Overview List Works

Actions are defined to run on the target object from different object levels, allowing you to add actions as recommendations for alert definitions that are configured for different base objects. The Actions overview is a list of actions available in your environment.

Where You Find the Actions Overview List

To view the available actions, in the menu, click **Alerts** and then in the left pane, click**Configuration > Actions**.

Table 2-28. Actions Overview Options

Option	Description
Filter options	Limits the list to actions matching the filter.
Action Name	Name of the action. Duplicate names indicate that the action name is provided by more than one adapter or has more than one associated object.
Action Type	Type of action that the action performs, either read or update. ■ Update actions make changes to the target objects. ■ Read actions retrieve data from the target objects.
Adapter Type	Name of the configured adapter that provides the action.
Resource Adapter Type	Adapter that provides the action.
Associated Object Types	Indicates the object level at which the action instance runs.
Recommendations	Indicates whether the action is used in at least one recommendation.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

■ `Set Memory for VM Power Off Allowed`

■ `Set CPU Count for VM Power Off Allowed`

■ `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

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 > Configuration > Alert Definitions**.

- Create, edit, and import recommendation definitions in **Alerts > Configuration > 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 2-29. Decreasing CPU Count and Memory Behavior Based On Options

Virtual Machine Power State	Power Off Allowed Selected	Results
On	Yes	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. If VMware Tools is not installed, the action powers off the virtual machine, decreases the CPU or memory, and powers the machine back on.
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 2‑30. 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	If VMware Tools is installed and running, the action shuts down the virtual machine, increases the CPU count, and powers the machine back on. If VMware Tools is not installed, the action powers off the virtual machine, increases the CPU count, and powers the machine back on.
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 2‑31. 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 ≤ 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.

Table 2-31. Increasing Memory Amount Behavior (continued)

Virtual Machine Power State	Memory Hot Plug Enabled	Hot Memory Limit	Power Off Allowed Selected	Results
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 thevRealize 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 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

- Workload Optimization Page

- Rightsizing

- Manage Optimization Schedules

- Workload Automation Policy Settings

- View DRS Summary

- Optimization Schedules

- Optimize Placement

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   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 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".

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 the **Vertical Ellipses** and then **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   If you need more information about the VMs which are included or excluded in the plan, click **Download Report** to see the optimization plan. You can review the reasons for incompatibilities and why some VMs were excluded from the plan.

7   Optional: If you want to know the total optimization potential of the move, assuming that there were no incompatibilities and all your VMs can be included in the optimization plan, click **Cancel**, and go to the Optimization Potential tab in the Workload Optimization page. Click **Calculate Optimization Potential** to see the total optimization potential of your data center.

8   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.

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>*.

Workload Optimization Page

Workload Optimization enables you to optimize virtual machines and storage across datastore clusters to reduce resource contention and maintain optimum system performance.

Where You Find Workload Optimization

From the Home screen, select **Workload Optimization** under Optimize Performance in the left pane. From the Quick Start screen, select **Workload Optimization** in the left-most column.

Workload Optimization Page Options

In the Workload Optimization page, you see a list of data centers in a carousel, listed under three categories:

- Critical

- Normal

- Unknown

After you select a data center, you see the **ALL DATACENTERS** button on the upper right. Click **ALL DATACENTERS** when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers.

Table 3-2. Workload Optimization Page Options

Option	Description
View:	Filter results to include data centers, custom data centers, vRA-managed custom data centers, or all three. (Option appears if you select **ALL DATACENTERS** on the upper right.)
Group By:	Filter results by criticality (most out of balance data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. (Option appears if you select **ALL DATACENTERS** on the upper right.)

Table 3-2. Workload Optimization Page Options (continued)

Option	Description
Sort By:	Options (Options appear if you select **ALL DATACENTERS** on the upper right): ■ Alarm clock graphic - list data centers/custom data centers by time remaining. ■ Dollar sign - list data centers/custom data centers by potential cost savings with capacity optimization. ■ Scales graphic - Optimized.
Select data center or **ADD NEW CUSTOM DATACENTER**	Options (Options appear if you select **ALL DATACENTERS** on the upper right): ■ Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object. ■ Select **ADD NEW CUSTOM DATACENTER** to display a screen that enables you to define a custom data center.

Data Center Options

After you select a data center from the carousel, you see the following information and options.

Note If you point your cursor to the lower right of a data center graphic, a tooltip may appear to let you know that the data center is using automated optimization.

Optimization Status Tab

Appears when you select a data center or custom data center from the top of the screen.

Table 3-3. The Optimization Recommendation Card

Option	Description
Status	■ Optimized - indicates that workloads are optimized based on the settings you entered in the neighboring Operational Intent window, with no tag violations based on the settings you entered in the Business Intent window. ■ Not Optimized - indicates that one of the following conditions is true: workloads are not optimized based on the settings you entered in the neighboring Operational Intent window AND/OR there are tag violations based on the settings you entered in the Business Intent window. In the event of tag violations, the offending tags are listed.
OPTIMIZE NOW	Runs optimizing actions based on the settings you entered in your Operational and Business Intent settings.

Table 3-3. The Optimization Recommendation Card (continued)

Option	Description
SCHEDULE	Displays a dialog box enabling you to schedule one or more optimization actions. If schedules are currently set for data center or custom data center optimization, a check mark appears next to the data center or custom data center name.
AUTOMATE	Continually seeks optimizing opportunities for data center or custom data center, based on the settings in the neighboring Operational Intent window or Business Intent windows. Scheduled optimizations are turned off while automatic optimization is on. Also, automated alerts are not operational when automatic optimization is on. Once you confirm automation, the system displays message, for example, 1) "Workload Optimization is looking for opportunities to automate," 2) "Your workloads are optimized according to your settings." or 3) "No eligible moves were found within the max number of compatibility checks allowed." **Note** To initiate Automation, you must have privileges for Environment -> Action -> Schedule Optimize Container.
TURN OFF AUTOMATION	Stops automatic optimization. Any scheduled optimizations come back online.

Note   Sometimes an optimizing action may be recommended, 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 recommended optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation can incur stress in the future, then consolidation is not recommended.

Table 3-4. The Operational Intent Card

Option	Description
Utilization Objective	Indicates the main attribute of your current automation policy settings. Values are moderate, consolidate, or balance.
Edit	Displays the Workload Automation Policy Settings, where you can adjust settings for optimization and cluster headroom.

Table 3-5. The Business Intent Card

Option	Description
Intent	Allows you to define zones of infrastructure within cluster boundaries.
Edit	Displays a workspace where you can select criteria for placement of VMs.

Table 3-6. Details for Are your clusters meeting your utilization objective?

Option	Description
Are your clusters meeting your utilization objective?	Displays a table which presents data in the following columns: ■ Name ■ CPU Workload ■ Memory Workload ■ DRS Settings ■ Migration Threshold ■ Violated Tags ■ VM Name Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. The violated tags shows which clusters or host groups are breaching the business intent. The VM Name column shows the name of the VMs and tag value due to which tag violation is happening. Provides the option to set the DRS automation level for individual objects.
VIEW DRS SUMMARY	Select a cluster in the list, then click this link to display a page containing metrics for DRS performance and cluster balance in the selected data center.
SET DRS AUTOMATION	Select a cluster in the list, then click this link to set the level of the DRS automation for the cluster. Note that clusters must be fully automated in order for workload optimization alerts to run actions set in the policies.

History Tab

Displays a graphical depiction of executed manual and automated optimizations for clusters in the selected data center or custom data center, based on parameters you provide.

Table 3-7. Details for History

Option	Description
Selected WLP process drop-down	The optimization action whose details you want to display.
Time duration drop-down	Last n hours - select the time parameter: last 6, 12, 24 hours or last 7 days.

Table 3-7. Details for History (continued)

Option	Description
Quick filter	choose a cluster name to search on.
Squares graphic	toggle between viewing processes in icon or circle form.
Circle	Toggle between viewing processes presented in a circle or on a straight line.
Back arrow - reset action.	Reset action.

If you point your cursor to a specific cluster as displayed on the screen, the details of the cluster appear in a tool tip. Click the note card icon on the lower right of the tool tip to go to the Details screen for the cluster. When displayed in the circle format, rings in the circle indicate how much CPU and how much memory was used at any given time. For example, if memory usage was higher than recommended based on your policy settings, the memory circle appears red.

Note the timeline across the bottom of the screen. When you choose parameters, for example, WLP process name, time parameter and cluster name, indicators appear along the timeline, showing when processes were initiated.

To zero in on a specific event, choose a process from the drop-down menu. You can also click points on the marker floating above the timeline, which causes a descriptive tool tip to appear, then double-click the 'Double-click to zoom' icon on the lower right.

If the event you choose includes an actual movement of VMs, you see a blue ball containing the number of VMs moved and showing the direction of the move and starting and ending clusters.

Optimization Potential Tab

When you run Workload Optimization, vRealize Operations Manager runs compatibility checks and excludes those VMs which have constraints, and only optimizes resources of those VMs which can be moved. If you want to see the total potential of your workload optimization, assuming that all VMs can be moved, click the **CALCULATE OPTIMIZATION POTENTIAL** button in the Optimization Potential tab. Optimization Potential disregards the underlying constraints and recommends moves before the compatibility checks. You can download the report to see more details.

If you want to see what can be realistically optimized, click **OPTIMIZE NOW** in the **Operation Status** tab. After you click **OPTIMIZE NOW**, you can download a report to review incompatibilities.

The optimization potential report helps you understand the difference between the optimization achievable when you run **OPTIMIZE NOW** and the total optimization potential.

See also Example: Run Workload Optimization

Rightsizing

Use this screen to alter the number of CPUs and amount of memory in oversized and undersized virtual machines.

Where You Find Rightsizing

From the Home screen, select **Rightsizing** under Optimize Capacity in the left pane.

Note Click on a data center graphic to display the details for the data center.

How Rightsizing Works

The Capacity Optimization, Reclaim, and Rightsizing features are tightly integrated functions that enable you to assess workload status and resource usage in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed. With this function, you can change CPU size and memory values for oversized and undersized virtual machines to achieve optimum system performance.

When you open the page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To identify possible oversized and undersized VMs in a data center, click its graphic. The area following refreshes to display details about the selected data center.

"Oversized VMs" displays the number of VMs determined to be oversized based on policies previously set. A chart details suggested reductions in the overall number of CPUs and GBs of memory and shows the percentage of total resources the reductions represent. Similarly, "Undersized VMs" indicates the number of VMs considered to be undersized, with a chart listing suggested increases in CPU and memory.

The table at the bottom of the page provides important information about the VMs. Table headings are Oversized VMs and Undersized VMs. VMs under each heading are grouped by cluster. Click the chevron to the left of a cluster name to list all the oversized or undersized VMs, respectively, in that cluster. You can check the box next to one or more VM names and click the **EXCLUDE VM(S)** button to prevent those VMs from being included in a resizing action. You can also select individual VMs to resize before clicking the **RESIZE VM(S)** button.

Run a Rightsize Action on Oversized VMs

Run the action as follows:

1   In the table headings, **Select** Oversized VMs.

2   **Select** the boxes next to VMs you want to exclude from the action, if any.

3   Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.

4   **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.

5   Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested reductions for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.

6   **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

Run a Rightsize Action on Undersized VMs

Run the action as follows:

1   In the table headings, **Select** Undersized VMs.

2   **Select** the boxes next to VMs you want to exclude from the action, if any.

3   Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.

4   **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.

5   Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested increases for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.

6   **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

Table 3-8. Rightsize Options

Option	Description
Select a data center.	Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.
ALL DATACENTERS | X	Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers.
View:	Filter results to include data centers, custom data centers, or both. Option appears when you select **ALL DATACENTERS** on the upper right.
Group BY:	Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select **ALL DATACENTERS** on the upper right.
Sort by:	Options (Options appear when you select **ALL DATACENTERS** on the upper right): ■ Alarm clock graphic - list data centers/custom data centers by time remaining. ■ Dollar sign - list data centers/custom data centers by potential cost savings. ■ Scales graphic - list data centers/custom data centers by level of optimization.

Table 3-8. Rightsize Options (continued)

Option	Description
Select data center or ADD NEW CUSTOM DATACENTER.	Options (Options appear when you select **ALL DATACENTERS** on the upper right): ■ Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. ■ Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that enables you to define a custom data center.
Oversized VMs display	Displays the number of VMs identified as oversized, with suggested reductions for vCPU and memory size.
Undersized VMs display	Displays the number of VMs identified as undersized, with suggested increases for vCPU and memory size.
Table of Oversized and Undersized VMs	Tabular representation of the Oversized and Undersized VMs in the selected data center. Click one of the headings - Oversized VMs or Undersized VMs - to refresh the table with data for that heading. The table lists the relevant VMs. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name. Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs. Once you select a VM or VMs, the dimmed options above the table become visible, as follows. **Exclude VM(s)**: the selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings. For Oversized VMs: ■ **RESIZE VM(s)**: the system displays a dialog box with suggestions for reducing vCPUs and memory. Click the edit icons to change resource size. For Underseized VMs:: ■ **RESIZE VM(s)**: the system displays a dialog box with suggestions for increasing vCPUs and memory. Click the edit icons to change resource size. **SHOW|HIDE EXCLUDED VMS**: toggle displays or hides the list of VMs you previously excluded. **INCLUDE VM(s)**: include the selected VMs in the actionable list.

Manage Optimization Schedules

Enables you to set up a regular schedule for optimizing a selected container.

Where You Find Manage Optimization Schedules

At the Workload Optimization screen, select **SCHEDULE** from the pane: Optimization Recommendation

Option	Description
Schedule Name	Meaningful name for the schedule
Time Zone	Choose the time zone for the action
Recurrence	Indicate how often you want the optimize action to run. Complex schedules can be defined, for example, select the Monthly option and choose to run the action on Tuesdays and every other Thursday, beginning on the fifth of the month.
Start on:	Day to start the optimization schedule.
Start at:	Time to start the optimization schedule.
Expire after:	Designate a set number of scheduled runs.
Expire on:	Designate an exact date for the actions to end.

See also Example: Schedule a Repeating Optimization Action

Workload Automation Policy Settings

Provides options for refining policy settings specifically for Workload Optimization.

Where You Find Workload Automation Settings

Access this screen through the Policies pages:

Select **Administration** from the menu, then select **Policies** from the left pane.

Click **Policy Library**, then click either the **Add New Policy** icon or the **Edit Selected Policy** icon. In the Add or Edit Monitoring policy workspace, on the left click **Workload Automation**.

Refer to Workload Automation Details .

View DRS Summary

The View DRS Summary page provides insight and perspective into the actions DRS is taking to balance a cluster. You can view DRS settings for the cluster and cluster balance metrics, and determine if recent vMotions are DRS- or user-initiated.

Where You Find the View DRS Summary Page

From the Home screen, select **Workload Optimization** under Optimize Performance in the left pane. Then select a cluster name in the Current Workloads pane. The dimmed View DRS Summary and Set DRS Automation links turn live. Click the link to display the DRS summary information.

Table 3-9. DRS Summary Values

Pane/fields	Value
<cluster name>	Name of the selected cluster
Automation Level	Enabled/Disabled. DRS is running or not.
Migration Threshold	Aggressive/Default/Moderate
Active Memory Used	False/ *nn%*
Cluster Balance	Shows the variations in the DRS cluster balance metric over time as DRS runs. The graph shows how DRS reacts to and clears any cluster imbalance each time it runs.
Cluster Imbalance	The range of potential imbalance values, as expressed in vCenter DRS metrics.
Total Imbalance	The level of imbalance in a cluster, as measured by vCenter DRS metrics.
Tolerable Threshold	The upper limit of what is tolerable in cluster imbalance. Designated by a green dotted line, this is a vCenter DRS metric.
VM Happiness	A bar graph summarizing the total happy and unhappy VMs in the cluster. For individual VMs, there is a presentation of performance metrics related to its happiness, such as %CPU ready time and memory swapped.
Happy VMs	Total of happy VMs are shown in green. Click in the green zone to show a list of these VMs in the Happy/Unhappy VMs pane to the right.
Unhappy VMs	Total of unhappy VMs is shown in red. To show a list of these VMs in the Happy/Unhappy VMs pane to the right, click in the red zone .
Happy/Unhappy VMs	Lists by name all the VMs in the zone you clicked in the VM Happiness pane.
VM Metrics	Shows the trend in VM happiness or unhappiness
Recent vMotions	The number of recent vMotions, plotted against time.
vMotion Details	Shows the number of DRS-initiated and user (non-DRS) initiated vMotions over time. You can choose which type you want to view.
Date/VM	Date of a given vMotion.
Source/Destination	Source and destination of moved VMs.
Type	DRS-initiated or user initiated.

Optimization Schedules

Use the Optimization Schedules page to edit or delete optimization schedules that you set up in the Manage Optimization Schedule Dialog Box at the Workload Optimization main screen.

Where You Find Optimization Schedules

- From the Home screen, select **Administration > Configuration > Optimization Schedules**.

- At the Workload Optimization Page page, select in the data center whose optimization schedule you want to edit or delete. Then click **SCHEDULE** in the Optimization Recommendation pane.

Table 3-10. Optimize Schedules Options

Option	Description
Edit icon	Select a schedule from the list, then click the **Edit** icon.
	The Manage Optimization Schedules appears, with the data for the selected schedule filled in.
Delete icon	Select a schedule from the list, then click the **Delete** icon.
	The selected schedule is deleted and does not run.

See also Example: Run Workload Optimization

Optimize Placement

A two-page dialog box that provides information about optimizing the workload of a selected container. When you run the optimization action, vRealize Operations Manager checks which of the VMs can be moved to a different cluster for better optimization of resources, based on the settings you entered in your Operational and Business Intent settings. You can download a report that provides information about the list of VMs that were included in, and excluded from, the move plan. The report provides reasons as to why some VMs were excluded from the plan.

First page: The current workload ("before," for example, CPU 105%) and projected results ("after," for example storage utilization 45%) for a possible optimizing action.

Second page: The exact moves planned for compute and storage resources.

Note It is possible that there is no optimization move plan. Review the report to see why vRealize Operations Manager could not provide a move plan.

Where You Find Optimize Placement

At the Workload Optimization screen, select OPTIMIZE NOW in the Optimization Recommendation pane.

Table 3-11. Optimize Clusters Options

Option	Description
Compare Cluster Balance	If you are satisfied with the before and after numbers (First page, above), click NEXT.
Review Optimization Moves	If you are satisfied with the moves planned (Second page, above), click BEGIN ACTION.
	Note Review the optimization plan report before you click BEGIN ACTION.
Download Report	The optimization plan report is in CSV format, and provides the following information:
	■ Summary of the optimization plan.
	■ Summary of the moves that make up the optimization plan.
	■ Issues related to the data center. Resolve these issues before proceeding with the optimization.
	■ Issues and incompatibilities applicable to specific VMs and their configurations. Resolve these issues, if applicable.
	■ Failed move attempts applicable to the specific VMs and their target destinations, as determined from the VM move plan. Resolve these issues and incompatibilities.

See also Example: Run Workload Optimization.

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 Policy Workspace to Create and Modify Operational Policies
- Define Monitoring Goals for vRealize Operations Manager Solutions

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 the policies in the priority order, as they appear in the priority column. 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 any active policy:

1   In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

    Note   The Reorder Policies option in enabled only if there are more than one active policies.

2   In the Reorder Policies window, select the policy and drag it up or down to change the priority.

3   Click **ok** to save the changes made to the priority.

The priority for the Default Policy is always designated with the letter D, and the other active policies are prioritized with numbers 1, 2, and so on. Policy with priority 1 indicates the highest priority. 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.

Table 4-1. Configurable Policy Rule Elements (continued)

Policy Rule Elements	Thresholds, Settings, Definitions
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:

1   In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

    Note   The Reorder Policies option in enabled only if there are more than one active policies.

2   In the Reorder Policies window, select the policy and drag it up or down to change the priority.

3   Click **ok** to save the changes made to the priority.

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.

Policies Library

The policies library displays the base settings, default policy, and other best practice policies that vRealize Operations Manager includes. You can use the policies library to create your own policies. The policies library includes all the configurable settings for the policy elements, such as workload, capacity and time remaining, and so on.

How the Policies Library Works

Use the options in policies library 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 or export a policy and reorder the policies.

Select a policy to display its details in the right pane. The right pane displays a high-level overview of all the details and options for that policy where these details are categorized in tabs. Expand each category to view all the related details.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for metrics and properties, alerts and symptoms, capacity, compliance, workload automation, and groups and objects. In this workspace, you can also apply the policy to objects and object groups. To update the policy associated with an object or object group, the role assigned to your user account must have the Manage Association permission enabled for policy management.

Where You Manage the Policies Library

To manage the policies library, in the menu, click **Administration**, and then in the left pane click **Policies**. The policies library appears and lists the policies available to use for your environment.

Table 4-2. Policy Library Tab Options

Option	Description
Toolbar	Use the toolbar selections to take action in the policies library. ■ Add. Create a policy from an existing policy. ■ Edit. Customize the policy so that you can override settings for vRealize Operations Manager to analyze and report data about the associated objects. ■ Delete. Remove a policy from the list. ■ 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 D, 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. ■ Reorder Policies. Change the priority of the active policies.
Policies library data grid	vRealize Operations Manager displays the high-level details for the policies. ■ Name. Name of the policy as it appears in the Add or Edit Policy workspace, 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. ■ Status: Indicates whether the policy is active or inactive.
Policies library > Right Pane	The right pane displays the name and description of the policy from which the settings are inherited, the policy priority, and the option to edit the policy. From the right pane, you can view the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created. ■ Metrics and Properties: Displays all the attribute types included in the policy. Attribute type includes, metrics properties, and super metrics. ■ Alerts and Symptoms: Displays all the alert and symptom definitions included in the policy. The Alert Definitions tabs display an overview of the alert definition, criticality, symptom, and state. The Symptoms Definitions tab displays an overview of the symptom name, criticality, and the metric name. ■ Capacity: Displays an overview of all the thresholds of the objects included in the policy. ■ Compliance: Displays the compliance thresholds inherited from the base policy or set while creating the policy. ■ Workload Automation: Displays the details of the workload optimized in your environment per your definition. ■ Groups and Objects: Displays the object or object groups associated with the selected policy and the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it is shown here.

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 objects or object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with objects or 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 objects or 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 or an object group, vRealize Operations Manager collects data from the objects 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 an individual object or 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 capacity 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 an individual object or groups of objects.

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.

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. 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 most of your objects.

The Default policy is marked with the letter D in the Priority column and can apply to any number of objects.

All the Default policies appear in the Default Policy group in the policies library, 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 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 in the policies library, these settings are called Base Settings. The Default policy inherits all 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** to see the policies provided with vRealize Operations Manager.

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 policies library, 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 policies library. All 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 configuration based policy set includes policies provided with vRealize Operations Manager that you use for specific settings on objects to report on your objects. This 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 most of your objects.

Using the Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the 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 objects or object groups in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated objects or object groups.

Prerequisites

Verify that objects or object 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 vRealize Operations Manager .

Procedure

1 In the menu, click **Administration**, and then in the left pane click **Policies**.

2 Click **Add** to add a policy, or select the policy and click **Edit Policy** to edit an existing policy.

You can add and edit policies 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 Create Policies workspace, assign a name to the policy, and enter the description.

Give the policy a meaningful name and description so that all users know the purpose of the policy.

4 From the **Inherit From** drop-down, select one or more policies to use as a baseline to define the settings for your new local policy.

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your new policy settings.

5 Click **Create Policy**.

The Create Policies workspace provides the options to customize your policy.

6 Click **Metrics and Properties**. In this 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.

a   Click **Save** and return to the create policies workspace.

7 Click **Alerts and Symptoms**. In this workspace, select the alert definitions and symptom definitions, and enable or disable them as required for your policy.

vRealize Operations Manager identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

a   Click **Save** and return to the create policies workspace.

8 Click **Capacity**. In this workspace, select and override the situational settings such as committed projects to calculate capacity, time remaining, and other detailed settings.

a   Click **Save** and return to the create policies workspace.

9 Click **Compliance**. In this workspace, set the compliance threshold required for your policy.

a   Click **Save** and return to the create policies workspace.

10 Click **Workload Automation**. In this workspace, select the optimization settings required for your policy.

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.

a   Click **Save** and return to the create policies workspace.

11 Click **Groups and Objects**. In this workspace, select one or more groups and objects to which the policy applies.

vRealize Operations Manager monitors the objects according to the settings in the policy that is applied to the object or 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 objects or object groups, vRealize Operations Manager does not assign the settings in that policy to any objects, and the policy is not active. For an object or an object group that dos not have a policy assigned, vRealize Operations Manager associates the object group with the Default 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.

a   Click **Save** and return to the create policies workspace.

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 objects or object groups.

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 objects or object groups in your environment. You can view the 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 the 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** and click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit the 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 objects or object groups.

To remove a policy from the list, select the policy, click the horizontal ellipse, and select **Delete**.

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 the Inherited 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 policy.

- Capacity 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.

- Compliance Details

  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 for the object types in your policy.

- 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 as per your definition.

- 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 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.

■ Groups and Objects details

You can assign your local policy to one or more objects or groups of objects to have vRealize Operations Manager analyze those objects according to the settings in your policy. You can 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** and click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. The name and description appear in the Create or Edit policy workspace.

Table 4-3. Name and Description Options in the Create or Edit Policy Workspace

Option	Description
Name	Name of the policy as it appears in the Create or Edit Policy screens, 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 must understand the relationship of the policy to one or more groups of objects.
Inherit From	The base policy that is 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. Select a base policy to inherit the policy settings as a starting point for your new policy.

Select the Inherited 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 policy.

In the policy content area, you can perform the following actions:

■ View the packages and elements for the inherited policy and additional policies that you selected to override the settings.

■ Compare the differences in settings highlighted between these policies.

■ Display object types.

To create a policy, select a base policy to inherit 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 to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the policy settings cards.

Click each card to display the inherited policy configuration, and your policy, and displays a preview of the selected policy settings. When you select one of the policy cards, you can view the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of enabled and disabled changes.

When you select the Groups and Objects card, 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, the workspace 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** and click **Add** to add a policy. In the Create policies workspace, add a name and description for the policy and from the **Inherit From** drop-down, select the base policy. The policy configuration, objects, and preview appear in cards below this drop-down.

Capacity 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 Capacity Workspace Works

When you turn on and configure the Capacity 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 edit a policy 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 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 changes these settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

Where You Set the Policy Capacity Settings

To set the capacity settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click the **Capacity** card. The capacity settings for host systems, virtual machines, and other object types that you select appear in the workspace.

You can also edit the capacity settings while working on the objects under the Environment Tab. In the **Capacity** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Capacity Setting**.

Table 4-4. Capacity Settings in the Create or Edit Policy Workspace

Option	Description
Select Object Type	Use the drop-down menu to select object types. Click the **All Filters** 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. Provide 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.
All Filters	When you select a filter, a list of the object types that you selected is displayed in the left pane with the threshold settings in the right pane.
Capacity settings for object types	Select an object to view the policy elements and settings for the object type so that you can have vRealize Operations Manager analyze the object type. You can view and modify the threshold settings for the following policy elements: ■ Workload ■ Time Remaining ■ Capacity Remaining ■ Maintenance Schedule ■ Allocation Model ■ Custom Profile ■ Capacity Buffer Click the lock icon on the left of each element to override the settings and change the thresholds for your policy.
Time Remaining Calculations	You can set the risk level for the time that is remaining when the forecasted total need of a metric reaches usable capacity. ■ 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 capacity setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The workload settings for the object types that you have 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 4-5. Policy Workload Element Settings in the Create or Edit Policies 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	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 capacity setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects. The time remaining settings for the object types that you have selected 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 4-6. Policy Time Remaining Element Settings in the Create or Edit Policies 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	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 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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The capacity remaining settings for the object types that you have selected 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 4-7. Policy Capacity Remaining Element Settings in the Create or Edit Policies 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	Allows you to set the percentage at which the capacity remaining alerts must be triggered.

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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The maintenance schedule settings for the object types that you have selected in the workspace 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 4-8. Policy Maintenance Schedule Element Settings in the Create or Edit Policies 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.

Policy Allocation Model Element

Allocation model defines how much CPU, memory, or disk space is allocated to objects in a cluster or datastore cluster. In the policy, you can turn on the Allocation Model element and configure the resource allocation for the objects.

How the Allocation Model Element Works

The Allocation Model element determines how calculates capacity when you allocate a specific amount of CPU, memory, and disk space resource to clusters or datastore clusters. You can specify the allocation ratio for either one, or all of the resource containers of the cluster. Unlike the demand model, the allocation model is used for capacity calculations only when you turn it on in the policy.

The allocation model element also affects the reclaimable resources for memory and storage in Reclaim page. When you turn on the Allocation Model element in the policy, the tabular representation of the VMs and snapshots in the selected data center from which resources can be reclaimed displays reclaimable memory and disk space based on the overcommit values.

Where You Override the Allocation Model Element

To view and override the policy workload analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The allocation model settings for the object types that you selected appear in the right pane.

Click the unlock icon next to Allocation Model to set the overcommit ratios.

Table 4-9. Policy Allocation Model Element Settings

Option	Description
Set overcommit ratio, to enable Allocation Model	Allows you to set the overcommit ratio for CPU, memory, or disk space. Select the check box next to the resource container you want to edit and change the overcommit ratio value.

Policy Custom Profile Element

The custom profile element lets you apply a custom profile which shows how many more of a specified object can fit in your environment depending on the available capacity and object configuration.

Where You Define the Custom Profiles

To define a custom profile, in the menu click **Administration**, and then in the left pane click **Configuration**. Click **Custom Profiles** and click the **Add** option to define a new custom profile.

Where You Select the Custom Profile Element

To view and override the policy Custom Profile analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The custom profile element for the object types that you selected in the workspace appear in the right pane. Click the lock icon to unlock the section and make changes.

Policy Capacity Buffer Element

The capacity buffer element lets you add buffer for capacity and cost calculation. For vCenter Server objects, you can add buffer to CPU, Memory, and Disk Space for the Demand and Allocation models. You can add capacity buffer to clusters and datastore clusters. The values that you define here affect the cluster cost calculation. The time remaining, capacity remaining, and recommended values are calculated based on the buffer. For WLP, capacity buffer is first considered and then the headroom that you have defined is considered.

Where You Define the Capacity Buffer

To view and override the policy Capacity Buffer analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The custom profile element for the object types that you selected in the workspace appear in the right pane. Click the lock icon to unlock the section and make changes.

How the Capacity Buffer Element Works

The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required. The value of the usable capacity reduces by the buffer amount that you specify here. The default buffer value is zero. If you are upgrading from a previous version of vRealize Operations Manager, the buffer values are carried forward to the new version.

The capacity buffer value that you specify for the Allocation model is considered only if you have enabled allocation model in the policy.

The following tables display the capacity buffer that you can define based on the vCenter Adapter object types:

Object Type	Valid Models for Capacity Buffer
CPU	Demand
	Allocation
Memory	Demand
	Allocation
Disk Space	Demand
	Allocation

Compliance Details

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 for the object types in your policy.

Where You Override the Policy Compliance

To view and override the policy compliance setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Compliance**

View the compliance thresholds 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. Compliance Settings in the Create or Edit Policies 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	Allows you to set the compliance score threshold based on the number of violations against those standards.

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 as 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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Workload Automation**.

Table 4-11. Workload Automation in the Create or Edit Policies Workspace

Option	Description
Workload Optimization	Select a goal for workload optimization.
	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.
	Select **Moderate** when you want to minimize the workload contention.
	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.
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.
	Note vSphere HA overhead is already included in useable capacity and this setting does not impact the HA overhead.
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.

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 **Configuration > 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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

You can also edit the metrics and properties while working on the objects under the Environment Tab. In the **Metrics** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Metrics Collection**.

Table 4-12. 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	Deselect the options in the **Attribute Type**, **State**, **KPI**, and **DT** drop-down menus, to narrow the list of attributes. ■ ✔ 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. 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.
Object Type	Filters the attributes list by object type.

Table 4-12. Metrics and Properties Options (continued)

Option	Description
Page Size	The number of attributes to list per page.
Attributes data grid	Display the attributes for a specific object type.

- 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 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 **Configuration > Alert Definitions**.

- To view the available symptom definitions, in the menu, click **Alerts** and then in the left pane click **Configuration > 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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The definitions appear in the workspace.

You can also edit the alert settings while working on the objects under the Environment Tab. In the **Alerts** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Alerts State**.

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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-13. Alert Definitions in the Create or Edit Policies 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	Deselect the options in the **Type** and **State** drop-down menus, to narrow the list of symptom definitions.
	Impact indicates the health, risk, and efficiency badges to which the alerts apply.
	Criticality indicates the information, critical, immediate, warning, or automatic criticality types to which the alert definition applies.
	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.
Object Type	Filters the alert definitions list by object type.
Page Size	The number of alert definitions to list per page.

Table 4-13. Alert Definitions in the Create or Edit Policies Workspace (continued)

Option	Description
Filter	Locates data in the alert definition list.
Alert Definitions data grid	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.

- Alert Definition. Meaningful name for the alert definition.
- Criticality. Indicates the criticality of the alert.
- Symptom. 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, disabled, 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 **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-14. Symptom Definitions in the Create or Edit Policies 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	Deselect the options in the **Type** and **State** drop-down menus, to narrow the list of symptom definitions. ■ ✔ 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. 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. State determines whether enabled, disabled, and inherited symptom definitions appear in the symptom definition list.
Object Type	Filters the symptom definitions list by object type
Page Size	The number of symptom definitions to list per page.

Table 4-14. Symptom Definitions in the Create or Edit Policies Workspace (continued)

Option	Description
Filter	Locate data in the symptom definition list.
Symptom Definitions data grid	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. ■ Symptom Definition. Symptom definition name as defined in the list of symptom definitions in the Content area. ■ Criticality. Indicates the criticality. ■ 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.

Groups and Objects details

You can assign your local policy to one or more objects or groups of objects to have vRealize Operations Manager analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

How the Groups and Objects Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more objects or groups of objects. 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 and Objects

To apply the policy to an object or groups of objects, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**.

Groups and Objects Options

To apply the policy to an object or groups of objects, select the check box for the groups or objects in the workspace.

You can then view the groups and objects associated with the policy. In the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**. Click the **Custom Groups** tab to apply the policy to one or more groups of objects. Click the **Objects** tabs to apply the policy to one or more objects.

For more information about how to create an object group, see the topic called **Custom Object Groups Workspace to Create a New Group**.

For more information about how to create a policy, see Policy Workspace in vRealize Operations Manager.

Define Monitoring Goals for vRealize Operations Manager Solutions

The Manage Solution configuration for the vSphere solution provides a set of questions for you to answer to help you define the default policy settings associated with your vCenter Adapter. You can create a policy for a management pack solution that you add to vRealize Operations Manager.

How Define Monitoring Goals Works in vRealize Operations Manager

The Manage Solution workspace includes an option to define monitoring goals for the solution. The selections you make determine the default policy settings that vRealize Operations Manager uses to analyze and monitor the objects associated with the solution.

For example, you might have a production environment that is composed of four separate production areas, each of which includes specific object groups. To monitor the objects in each production area, you must set the default policy settings according to the monitoring requirements for each area. You can have vRealize Operations Manager set the default settings based on your infrastructure or virtual machines, alert you on individual objects or object groups, and so on.

Where You Define the Monitoring Goals for a Solution

To define the monitoring goals for a solution and establish the default settings for monitoring goals in the default policy, in the menu, click **Administration**, and then in the left pane, click **Solutions > Configuration**, and select a solution. Click **Configure**, and click **Define Monitoring Goals**. In the Define Monitoring Goals dialog box that appears, select answers to the questions about your objects, alerts, memory capacity, and compliance settings according to the *vSphere Hardening Guide*.

When you select an option, vRealize Operations Manager saves your setting. If you display the Define Monitoring Goals dialog box later, and the user interface did not appear to retain your selection, the selection is still active. As a double-check, select the option again, and click **Save**.

To adjust advanced settings of the policy, in the menu, click **Administration**, and then in the left pane, click **Policies**.

Table 4-15. Define Monitoring Goals Questions

Option	Description
Which objects do you want to be alerted on in your environment?	Select the type of objects to receive alerts. You can have vRealize Operations Manager alert on all infrastructure objects except for virtual machines, only virtual machines, or all.
Which type of alerts do you want to enable?	You can enable vRealize Operations Manager to trigger Health, Risk, and Efficiency alerts on your objects.
Configure Memory Capacity based on?	Set the memory capacity model based on the type of environment to monitor. For example, to monitor a production environment, select the **vSphere Default** model to use moderate settings to ensure performance. Use **Most Aggressive** for test and development environments. Use **Most Conservative** to use all allocated memory for capacity calculations.
Enable *vSphere Hardening Guide* Alerts?	Use the *vSphere Hardening Guide* to continuously and securely assess and operate your vSphere objects. When you enable these alerts, vRealize Operations Manager assesses your objects against the *vSphere Hardening Guide* rules. vSphere 6.0 objects are assessed against vSphere 6.0 hardening rules, and vSphere 5.5 objects are assessed against vSphere 5.5 hardening rules.
Learn More links	To display more information about a monitoring goal selection, click **Learn More**.

You can find the *vSphere Hardening Guides* at http://www.vmware.com/security/hardening-guides.html.

Configuring Compliance

<div style="text-align: right; font-size: 3em;">5</div>

You can set compliance on your objects to meet the defined standards and determine the compliance of your objects against the configuration standards.

This chapter includes the following topics:

- What Are Compliance Benchmarks
- How To Configure Compliance Benchmarks

What Are Compliance Benchmarks

Compliance benchmarks display score cards that help you proactively detect compliance problems in vRealize Operations Manager. The compliance benchmarks are measured against a set of standard rules, regulatory best practices, or custom alert definitions.

How Compliance Benchmarks Work

All the compliance standards in vRealize Operations Manager, including any standards that you define, are based on alert definitions. Only alert definitions of the Compliance subtype are counted. Custom score cards can monitor user-defined alerts.

In previous releases of vRealize Operations Manager, you had to modify the current default policy to monitor compliance against a set of standard rules, regulatory best practices, or custom alert definitions. In the current release, you can manage all compliance related tasks from the **Home > Troubleshoot > Compliance** page. When you configure a benchmark, you select an applicable policy. vRealize Operations Manager then enables the appropriate alert definitions in the policy to measure compliance.

The compliance assessment is based on the environment where your objects are deployed. You can monitor objects that are deployed in your VMware Self-Managed Cloud (SDDC) environment, including DC and Edge environments, and your VMware Managed Cloud (VMC SDDC) environment. Compliance benchmarks on VMC SDDC are applicable only on client VMs that you have deployed in the VMware Managed Cloud environment.

vRealize Operations Manager Compliance Benchmark Types

VMware SDDC Benchmarks

Displays score cards based on alerts which are measured against the latest hardening guides:

- vSphere Security Configuration Guide

- vSAN Security Configuration Guide

- NSX Security Configuration Guide

Displays benchmarks for and in the SDDC and VMC SDDC tabs.

> **Note**   vSphere 6.7 Update 1 Security Configuration Guide no longer contains risk profiles. For more information, see blogs.vmware.com.

Custom Benchmarks

Displays benchmarks that you define. Use compliance alerts from vSphere and regulatory management packs, or define your own alerts to monitor. You can define up to five custom score cards. You can import custom score cards from other instances of vRealize Operations Manager.

Regulatory Benchmarks

Displays benchmarks for industry standard regulatory compliance requirements. You can install compliance packs for the following regulatory standards:

- Health Insurance Portability and Accountability Act (HIPAA)

- Payment Card Industry Data Security Standard (PCI DSS) compliance standards

- CIS Security Standards

- Defense Information Systems Agency (DISA) Security Standards

- The Federal Information Security Management Act (FISMA) Security Standards

- International Organization for Standardization (ISO) Security Standards

For instructions on installing these compliance packs, see Install a Regulatory Benchmark.

Compliance Score Cards

The compliance page in vRealize Operations Manager displays score cards for each type of benchmark. A score card is a compliance visualization term.

What is a Compliance Score Card

Score cards in the Compliance landing page display the number of non-compliant objects, and the total number of objects affected by each hardening guide as well as the compliance score which is counted as the ratio of compliant objects to total number of objects assessed by the given benchmark, represented in percentage. In addition, you can see the breakdown of the total number of objects that are compliant and non-compliant. You can click on a score card to see more details, including alerts that were triggered based on the compliance standards.

The compliance score card of an object is counted as the smallest rounded off integer (100*(total number of symptoms triggered on an object/total number of symptoms)).

The compliance score for the object is based on the most critical of the violated standards. The score card displays 100 when all objects are compliant. When an object is non-compliant, the number of non-compliant symptoms are displayed in red and the total number of symptoms in grey.

Where You Find Compliance Score Cards

You can view score cards for each of the different types of benchmarks in the **Home > Troubleshoot > Compliance** page.

You can view score cards for objects in the **Environment > Object > Compliance** tab.

Compliance Page

In the **Home > Troubleshoot > Compliance** summary page, vRealize Operations Manager monitors compliance for SDDC and VMC SDDC objects. You can switch between the tabs to view the benchmarks for your on-premise deployment and cloud environments.

In each of these tabs, vRealize Operations Manager displays compliance score cards in the following sections:

- VMware SDDC Benchmarks

- Custom Benchmarks

- Regulatory Benchmarks

Compliance Tab

In the **Environment > Object > Compliance** tab, vRealize Operations Manager displays score cards for the benchmarks that include the current objects in their calculations, based on the alert definitions and policies associated with that benchmark. The score cards display the total number of rules and the number non-compliant (violated) rules based on symptoms for each hardening guide.

Score Cards in the Compliance Page

In the **Home > Troubleshoot > Compliance** page, you can view scores for benchmarks that you have enabled. Click a score card to view more information.

Table 5-1. Compliance Page Score Card Options

Item	Description
Score card for the configured hardening guides, custom benchmark and management packs	Displays the compliance score, total compliant and non-compliant objects for the compliance standards you have configured.
Object Breakdown	Displays the number of compliant and non-compliant objects for the following types of objects: ■ vCenter ■ ESXi Host ■ Virtual Machine ■ Distributed Port Group ■ Distributed Virtual Switch ■ vSAN Cache Disk ■ vSAN Capacity Disk ■ vSAN Cluster ■ NSX-T Manager ■ NSX-V EDGE ■ NSX-V Logical Router ■ NSX-V Manager ■ NSX-V Routing Edge Service
Compliance Alert List	A list of alerts, grouped by time by default. You can either remove the grouping of the alerts, or group by criticality, definition, and object type. The alerts which caused the compliance violation are displayed in a table. You can sort the table by the following columns: ■ Alert ID ■ Criticality ■ Alert ■ Triggered On ■ Updated On Select an alert from the table and click **Actions** to perform tasks such as canceling the alert, suspending alert, and taking ownership of the alert. Click an alert to view more details. The **Environment > Object > Alert** tab opens.

Compliance Alerts

You use the compliance score card as an investigative tool when you evaluate the state of objects in your environment, or when you research the root cause of a problem. If the score card indicates a problem, you can view the alerts to see details about the violation. Violated rules are based on the symptoms defined in the compliance alert.

The compliance alerts, which have the subtype named Compliance, include one or more symptoms that represent the compliance rules. Compliance alerts that are triggered appear on the **Environment > Object > Compliance** tab as violations to the standard, and the triggered symptoms appear as violated rules. The rules are the alert symptoms, and the symptom configuration identifies the incorrect value or configuration. If a rule symptom is triggered for any of the alerts in the standard, the triggered rule violates the standard and affects the score that appears on the **Environment > Object > Compliance** tab.

Table 5-2. Compliance Tab Alert Display

Item	Description
Score card for the configured hardening guides	Displays the score card value, total number of rules, and number of non-compliance rules for the compliance standards you have configured.
Active Compliance Alerts	If you click the score card, the rules for the score card appear. When a symptom is triggered, the rule is considered to be violated. View the list of rules in the following tabs: ■ Violated Rules. Displays only the triggered symptoms. Click a symptom to view more information. ■ All Rules. Displays triggered and untriggered symptoms.

How To Configure Compliance Benchmarks

Configure VMware SDDC, custom, and regulatory benchmarks from the Compliance page. Unlike previous releases, you can now enable alert definitions in one of the active policies, from the Compliance page directly.

Enable VMware SDDC Benchmarks

You can enable the VMware SDDC Benchmark to monitor objects for violation of vSphere Security Configuration Guide, vSAN Security Configuration Guide, NSX Security Configuration Guide (SDDC only). The score cards in the VMware SDDC Benchmark warn you when compliance alerts trigger on your vCenter Server instance, NSX-V objects, NSX-T objects, vSAN objects, ESXi hosts, virtual machines, distributed port groups, or distributed virtual switches.

Procedure

1   Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.

2   To enable the Security Configuration Guides, select either the SDDC or the VMC SDDC tab depending on the environment where your objects are present.

3   In the VMware SDDC Benchmarks section, click **Enable** under the vSphere Security
    Configuration Guide or vSAN Security Configuration Guide pane.

> **Note**   To enable the NSX Security Configuration guide, you must first install the NSX for
> vSphere, or the NSX-T solution. For more details, see Adding Solutions .

The **Enable Policies** dialog box opens.

4   Select the policy that you want to modify. When there are child policies, you can select a
    child policy and unselect a parent policy. vRealize Operations Manager modifies the selected
    policy and enables the alert definitions associated with the current scorecard.

5   Click **Enable** to confirm your selection.

Results

vRealize Operations Manager starts to assess the objects based on the policy that you selected.
To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

Create a New Custom Benchmark

You can create a custom compliance benchmark to ensure that objects comply with compliance
alerts available in vvRealize Operations Manager, or custom compliance alert definitions. When a
compliance alert is triggered on your vCenter instance, hosts, virtual machines, distributed port
groups, or distributed switches, you investigate the compliance violation. You can add up to five
custom compliance score cards.

Prerequisites

To create a custom benchmark based on industry standard regulatory compliance requirements,
you must first download and install the compliance management packs.

Procedure

1   Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.

2   To create a custom benchmark, first select either the SDDC or the VMC SDDC tab depending
    on where your objects are present.

3   In the Custom Benchmarks section, click **Add Custom Compliance**.

    The **Add Custom Compliance** dialog box opens.

4   Select **Create a New Custom Benchmark**.

    a   In the Name and Description step, provide a name and description for the custom
        benchmark and click **Next**.

    b   In the Alert Definitions step, select the compliance alerts that you want to add to this
        custom compliance benchmark and click **Next**.

    c   In the Policies step, select the policies to enable compliance and click **Finish**.

Results

The custom compliance which monitors alert definitions that you selected is available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit**.

Import or Export a Custom Benchmark

You can export custom benchmarks from any vRealize Operations Manager instance and import it to another instance. Reusing custom benchmarks saves you time and effort. You can modify an imported custom benchmark. Exported files are in the XML format. The XML file contains information about alert groups, alerts, and filters.

Prerequisites

You must first export a XML file with the custom benchmarks from another instance of vRealize Operations Manager before importing the XML file to another instance.

Procedure

1   Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.

2   To import a custom benchmark, select either the SDDC or the VMC SDDC tab depending on where your objects are present.

3   In the Custom Benchmarks section, click **Add Custom Compliance**.

The **Add Custom Compliance** dialog box opens.

4   Select **Import An Existing Custom Benchmark**.

    a   In the Import Compliance Score card dialog box, select the Score card Definition XML file from your local computer. If the XML file contains cloned alerts from the vRealize Operations Manager instance that was used to export the file, the cloned alerts are also imported.

    b   vRealize Operations Manager displays a message to indicate if the XML file was successfully imported.

    c   If you see a message which indicates that there is a conflict between the data in the XML file and the custom benchmarks already defined, make a selection on how to handle a conflict.

    d   Click **Done**.

5   To export an existing custom benchmark, click the score card to select the benchmark and select **Export** from the **Actions** menu.

Results

The imported compliance benchmarks are available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit** from the **Actions** menu after clicking the score card.

Install a Regulatory Benchmark

To enforce and report on the compliance of your vSphere objects, you activate the compliance pack that contains the policies for regulatory standards. Then, you select the policy to enable the appropriate regulatory alerts for your virtual machines.

Procedure

1 Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance**.

   The compliance packs for the regulatory standards are displayed under the Regulatory Benchmark section.

2 To install any regulatory benchmark, click **Activate From Repository** on the required compliance pack.

   You are redirected to the **Native Management Packs** page.

3 Navigate to the required compliance pack and click **Activate** to complete the installation.

4 To enable the compliance pack policies, navigate to the **Compliance** homepage and click **Enable** on the installed compliance pack.

   The **Enable Policies** window opens.

5 Select the policies that you want to enable and click **Enable** to complete the process.

Results

vRealize Operations Manager starts to assess the objects based on the regulatory benchmark that you installed.

Configuring Super Metrics

<div style="text-align: right; font-size: 4em; color: #ccc">6</div>

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 6-1. Designing a Super Metric Checklist

☐ 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.
☐ 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.
☐ 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 6-1. Designing a Super Metric Checklist (continued)

☐ 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.
☐ 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
- Super Metrics Tab

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.

The **Manage Super Metric** wizard opens.

3   Enter a meaningful name for the super metric such as `Worst VM CPU Usage (%)` in the
Name text box.

> **Note**   It is important that you have an intuitive name as it appears in dashboards, alerts, and
> reports. For meaningful names, always use space between words so that it is easier to read.
> Use title case for consistency with the out of the box metrics and add the unit at the end.

4   Provide a brief summary of the super metric in the **Description** text box.

> **Note**   Information regarding the super metric, like why it was created and by whom can
> provide clarity and help you track your super metrics with ease.

5   Select the unit of the super metrics from the **Unit** drop-down and click **Next**.

> **Note**   The super metrics unit configured here can be changed in the metrics charts, widgets,
> and views.

The Create a formula screen appears.

6   Create the formula for the super metric.

For example, to add a super metric that captures the average CPU usage across all virtual
machines in a cluster, perform the following steps.

a   Select the function or operator. This selection helps combine the metric expression with
operators and/or functions. In the super metric editor, enter `avg` and select the **avg**
function.

You can manually enter functions, operators, objects, object types, metrics, metrics types,
property, and properties types in the text box and use the suggestive text to complete
your super metric formula.

Alternatively, select the function or operator from the **Functions** and **Operators** drop-
down menus.

b   To create a metric expression, enter `Virtual` and select **Virtual Machine** from the object
type list.

c   Add the metric type, enter `usage`, and select the **CPU|Usage (%)** metric from the metric type list.

Note   The expression ends with depth=1 by default. If the expression ends with depth=1, that means that the metric is assigned to an object that is one level above virtual machines in the relationship chain. However, since this super metric is for a cluster which is two levels above virtual machine in the relationship chain, change the depth to 2.

The depth can also be negative, this happens when you need to aggregate the parents of a child object. For example, when aggregating all the VMs in a datastore, the metric expression ends with depth=-1, because VM is a parent object of datastore. But, if you want to aggregate all the VMs at a Datastore Cluster level, you need to implement 2 super metrics. You cannot directly aggregate from VM to Datastore Cluster, because both are parents of a datastore. For a super metric to be valid, depth cannot be 0 (-1+1=0). Hence, you need to create the first super metric (with depth=-1) for the aggregate at the datastore level, and then build the second super metric based on the first (with depth = 1).

The metric expression is created.

d   To calculate the average CPU usage of powered on virtual machines in a cluster, you can add the `where` clause. Enter `where=""`.

Note   The `where` clause cannot point to another object, but can point to a different metric in the same object. For example, you cannot count the number of VMs in a cluster with the CPU contention metric > SLA of that cluster. The phrase "SLA of that cluster " belongs to the cluster object, and not to the VM object. The right operand must also be a number and cannot be another super metric or variable. The where clause cannot be combined using AND, OR, NOT, which means you cannot have `where="VM CPU>4 and VM RAM>16"` in your super metric formula.

e   Position the pointer between the quotation marks, enter **Virtual**, and select the **Virtual Machine** object type and the **System|Powered ON** metric type.

f   To add the numeric value for the metric, enter **==1**.

g   To view hints and suggestions, click **ctrl+space** and select the adapter type, objects, object types, metrics, metrics types, property, and properties types to build your super metric formula.

h   Click the **This object** icon.

If the **This object** icon is selected during the creation of a metric expression, it means that the metric expression is associated to the object for which the super metric is created.

7   You can also use the **Legacy** template to create a super metric formula without the suggestive text.

To view the super metric formula in a human-readable format, click the **Show Formula Description** icon. If the formula syntax is wrong, an error message appears.

Note   If you are using Internet Explorer, you are automatically directed to the legacy template.

8   Verify that the super metric formula has been created correctly.

a   Expand the **Preview** section.

b   In the **Objects** text box, enter and select a **Cluster**.

A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

c   Click the **Snapshots** icon.

You can save a snapshot, or download the metric chart in a `.csv` format.

d   Click the **Monitoring Objects** icon.

If enabled, only the objects that are being monitored are used in the formula calculation.

e   Click **Next**.

The Assign to Object Types screen appears.

9   Associate the super metric with an object type. vRealize Operations Manager calculates the super metric for the target objects and displays it as a metric for the object type.

a   In the **Assign to an Object Type** text box, enter **Cluster** and select the **Cluster Compute Resource** object type.

After one collection cycle, the super metric appears on each instance of the specified object type. For example, if you define a super metric to calculate the average CPU usage across all virtual machines and assign it to the cluster object type, the super metric appears as a super metric on each cluster.

b   Click **Next**.

The Enable in a Policy screen appears.

10   Enable the super metric in a policy, wait for at least one collection cycle till the super metric begins collecting and processing data, and then review your super metric on the **All Metrics** tab.

a   In the **Enable in a Policy** section, you can view the policies related to the object types you assigned your super metric to. Select the policy in which you want to enable the super metric. For example, select the **Default Policy** for Cluster.

11 Click **Finish**.

You can now view the super metric you created and the associated object type and policy on the **Super Metrics** page.

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=metric_group|my_metric} > 0)`.

For example: `count(${objecttype = ExampleAdapter, adaptertype = ExampleObject, metric = ExampleGroup|Rating, depth=2, where =($value==1})`

IsFresh Function

Use the `isFresh` function in the `where` clause to check if the last value of the metrics is fresh or not.

For every metric published in vRealize Operations Manager, the point with the latest publishing time is called as the last point of that metric. The value of that metric's last point is called the last value of that metric. A metric's last point is considered fresh when the time elapsed after the metric's last point is lesser than the estimated publishing interval of that metric.

The `isFresh` function returns true if the last value of the metrics is fresh. For example, in the following scenarios, the function:

- `${this, metric=a|b, where=($value.isFresh())}`, returns the last value of the metric a|b if the last value is fresh.

- `${this, metric=a|b, where=($value == 7 && $value.isFresh())}`, returns the last value of the metric a|b if it is equal to seven and is fresh.

- `${this, metric=a|b, where=(${metric=c|d} == 7 && ${metric=c|d}.isFresh())}`, returns the last value of the metric a|b only if the last value of the metric c|d is equal to seven and is fresh.

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(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)/
(max(${adaptertype=VMWARE, objecttype=HostSystem, attribute=cpu|demand|
active_longterm_load, depth=5, where=($value>=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(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=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.

Super Metrics Tab

A super metric is a mathematical formula that contains a combination of one or more metrics for one or more objects. With super metrics you can assess information more quickly when you are observing fewer metrics.

Where You Configure Super Metrics

Click **Administration** and in the left pane click **Configuration > Super Metrics**.

Table 6-2. Configuration Options for Super Metrics

Option	Description
Toolbar	Use the toolbar selections to manage super metric options. ■ Add New Super Metric. Starts the Manage Super Metric workspace. See Manage Super Metric Workspace. ■ Edit Selected Super Metric. Starts the Manage Super Metric workspace. ■ Clone Selected Super Metric. Duplicates the super metric. Edit the clone or associate it with a different object type. ■ Delete Selected Super Metric. ■ Export Selected Super Metric. Exports a super metric to use in another vRealize Operations Manager instance. See Exporting and Importing a Super Metric. ■ Import Super Metric. Imports a super metric to this vRealize Operations Manager instance. See Exporting and Importing a Super Metric.
Super Metrics list	Configured super metrics listed by name and formula description.
Policies Tab	Policies in which the super metric attribute is enabled for collection. When enabled in a policy, vRealize Operations Manager collects super metrics from the objects associated with the policy. See Metrics and Properties Details.
Object Types Tab	Object types for the super metric display. vRealize Operations Manager calculates the super metric for the objects associated with the object type and displays the value with the object type. Use the toolbar selections to add or delete an object type association.

Manage Super Metric Workspace

You use the Manage Super Metric workspace to create or edit a super metric. The toolbar helps you to build the mathematical formula with the objects and metrics you select.

Where You Configure Super Metrics

On the menu, click **Administration** and in the left pane click **Configuration > Super Metrics**.

Table 6-3. Super Metrics Workspace Options

Option	Description
Super Metric	Use the toolbar selections to build and display your super metric formula. ■ Functions. Mathematical functions that operate on a single object or group of objects. See Super Metric Functions and Operators. ■ Operators. Mathematical symbols to enclose or insert between functions. See Enhancing Your Super Metrics. ■ This Object. Assigns the super metric to the object selected in the Object pane and displays `this` in the formula instead of a long description for the object. ■ Show Formula Description. Shows the formula in a textual format. ■ Visualize Super Metric. Shows the super metric in a graph. Look at the graph so that you can verify that vRealize Operations Manager is calculating the super metric for the target objects that you selected. ■ Name. The name you give to the super metric.
Objects Pane	Displays the list of objects collecting metrics. Use this list to select the object with the metrics to measure. If an object type is selected, only objects of the selected type are listed. Column headings help you to identify the object.
Object Types Pane	Use this list to select the object type with the metrics to measure. The object type selection affects the list of objects, metrics, and attribute types displayed. ■ Adapter Type. Shows the object types for the adapter selected. ■ Filter. Shows the object types with the filter words.
Metrics Pane	Displays the list of available metrics for the object or object type selection. Use this list to select the metrics to add to the formula.
Attribute Types Pane	Displays the list of attribute types for the object or object type selection. Use this list to select the metrics for the attribute type to add to the formula.

Super Metric Functions and Operators

vRealize Operations Manager includes functions and operators that you can use in super metric formulas. The functions are either looping functions or single functions.

Looping Functions

Looping functions work on more than one value.

Table 6-4. Looping Functions

Function	Description
avg	Average of the collected values.
combine	Combines all the values of the metrics of the included objects in a single metric timeline.
count	Number of values collected.
max	Maximum value of the collected values.
min	Minimum value of the collected values.
sum	Total of the collected values.

Note vRealize Operations Manager 5.x included two sum functions: sum (expr) and sumN (expr, depth). vRealize Operations Manager 6.x includes one sum function: sum (expr). Depth is set at depth=1 by default. For more information about setting depth, refer to Create a Super Metric.

Looping Function Arguments

The looping function returns an attribute or metric value for an object or object type. An attribute is metadata that describes the metric for the adapter to collect from the object. A metric is an instance of an attribute. The argument syntax defines the desired result.

For example, CPU usage is an attribute of a virtual machine object. If a virtual machine has multiple CPUs, the CPU usage for each CPU is a metric instance. If a virtual machine has one CPU, then the function for the attribute or the metric return the same result.

Table 6-5. Looping Function Formats

Argument syntax example	Description
funct(${this, metric =*a|b:optional_instance|c*})	Returns a single data point of a particular metric for the object to which the super metric is assigned. This super metric does not take values from the children or parents of the object.
funct(${this, attribute=*a|b:optional_instance|c*})	Returns a set of data points for attributes of the object to which the super metric is assigned. This super metric does not take values from the child or parent of the object.
funct(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={id1=*val1*id2=*val2*,...}, metric=*a|b:instance|c*})	Returns a single data point of a particular metric for the *resname* specified in the argument. This super metric does not take values from the children or parents of the object.
funct(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={id1=*val1*, id2=*val2*,...}, attribute=*a|b:optional_instance|c*})	Returns a set of data points. This function iterates attributes of the *resname* specified in the argument. This super metric does not take values from the child or parent of the object.

Table 6-5. Looping Function Formats (continued)

Argument syntax example	Description
funct(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, metric=*a*/*b*:*optional_instance*/*c*})	Returns a set of data points. This function iterates metrics of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects, where *depth* describes the object location in the relationship chain.
	For example, a typical relationship chain includes a data center, cluster, host, and virtual machines. The data center is at the top and the virtual machines at the bottom. If the super metric is assigned to the cluster and the function definition includes depth = 2, the super metric takes values from the virtual machines. If the function definition includes depth = -1, the super metric takes values from the data center.
funct(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, attribute=*a*/*b*:*optional_instance*/*c*})	Returns a set of data points. This function iterates attributes of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects.

For example, `avg(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|usage_average, depth=1})` averages the value of all metric instances with the cpu|usage_average attribute for all objects of type `VirtualMachine` that the vCenter adapter finds. vRealize Operations Manager searches for objects one level below the object type where you assign the super metric.

Single Functions

Single functions work on only a single value or a single pair of values.

Table 6-6. Single Functions

Function	Format	Description
abs	abs(x)	Absolute value of x. x can be any floating point number.
acos	acos(x)	Arccosine of x.
asin	asin(x)	Arcsine of x.
atan	atan(x)	Arctangent of x.
ceil	ceil(x)	The smallest integer that is greater than or equal to x.
cos	cos(x)	Cosine of x.
cosh	cosh(x)	Hyperbolic cosine of x.
exp	exp(x)	e raised to the power of x.
floor	floor(x)	The largest integer that is less than or equal to x.
log	log(x)	Natural logarithm (base *x*) of x.
log10	log10(x)	Common logarithm (base 10) of x.
pow	pow(x,y)	Raises x to the y power.

Table 6-6. Single Functions (continued)

Function	Format	Description
rand	rand()	Generates a pseudo random floating number greater than or equal to 0.0 and less than 1.0.
sin	sin(x)	Sine of x.
sinh	sinh(x)	Hyperbolic sine of x.
sqrt	sqrt(x)	Square root of x.
tan	tan(x)	Tangent of x.
tanh	tanh(x)	Hyperbolic tangent of x.

Operators

Operators are mathematical symbols and text to enclose or insert between functions.

Table 6-7. Numeric Operators

Operators	Description
+	Plus
-	Subtract
*	Multiply
/	Divide
%	Modulo
==	Equal
!=	Not equal
<	Less than
<=	Less than, or equal
>	Greater than
>=	Greater than, or equal
||	Or
&&	And
!	Not
? :	Ternary operator. If/then/else For example: **conditional_expression ? expression_if_condition_is_true : expression_if_condition_is_false** For more information about ternary operators, see Enhancing Your Super Metrics.

Table 6-7. Numeric Operators (continued)

Operators	Description
( )	Parentheses
[ ]	Use in an array of expressions
[x, y, z]	An array containing x, y, z. For example, min([x, y, z])

Table 6-8. String Operators

String Operators	Description
equals	Returns true if metric/property string value is equal to specified string.
contains	Returns true if metric/property string value contains specified string.
startsWith	Returns true if metric/property string value starts with the specified prefix.
endsWith	Returns true if metric/property string value ends with the specified suffix.
!equals	Returns true if metric/property string value is not equal to specified string.
!contains	Returns true if metric/property string value does not contain specified string.
!startsWith	Returns true if metric/property string value does not start with the specified prefix.
!endsWith	Returns true if metric/property string value does not end with the specified suffix.

Note String operators are valid in 'where' condition only. For example: **${this, metric=summary|runtime|isIdle, where = "System Properties| resource_kind_type !contains GENERAL"}**

Configuring Objects

7

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, that are the central components of management packs. You can customize adapter instances for your virtual environment using cloud accounts and other accounts. vRealize Operations Manager uses cloud accounts to manage the communication and integration with other products, applications, and functions.

■ Cloud Accounts - You can configure cloud adapter instances and collect data from cloud solutions that are already installed in your cloud environment from the cloud accounts page.

■ Other Accounts - You can view and configure native management packs and other solutions that are already installed and configure adapter instances from the other accounts page.

■ Repository - You can activate or deactivate native management packs and add or upgrade other management packs from the Repository page.

The screenshot displays the list of available solutions in vRealize Operations Manager . You must first Activate the solution before adding and configuring the accounts.

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.

Workload Management Inventory Objects

vRealize Operations Manager discovers the following workload management objects and their child objects using the vCenter adapter:

■ Tanzu Kubernetes cluster

■ vSphere Pods

- Namespace

A cluster with Kubernetes enabled, running on vSphere, is called a Supervisor Cluster. In the vRealize Operations Manager inventory, the summary tab of the Supervisor Cluster indicates that it has workload management enabled. The Supervisor Cluster contains specific objects that enable the capability to run Kubernetes workloads within ESXi. vRealize Operations Manager collects metrics and data for the Supervisor Cluster. Supervisor Clusters contain Namespaces, which are resource pools that have dedicated memory, CPU, and storage.

Namespaces contain virtual machines with k8s enabled. They are called k8s control VMs. These VMs are managed by vSphere. Therefore, you cannot take action on these VMs from within vRealize Operations Manager.

DevOps engineers can run workloads on containers running inside vSphere Pods. They can create Tanzu k8s cluster inside a Namespace. A vSphere Pod is a VM with a small footprint that runs one or more Linux containers. It is the equivalent of a k8s pod. A Tanzu Kubernetes cluster is a full distribution of the open-source Kubernetes container orchestration software that is packaged, signed, and supported by VMware.

To understand the vSphere Tanzu Kubernetes architecture, see *Configuring and Managing vSphere with Kubernetes* in the vSphere documentation.

Workload management objects are excluded from the following workflows:

- Compliance

- Reclaim

- Rightsizing

- Workload optimization

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 Chapter 1 Connecting vRealize Operations Manager to Data Sources.

Verify that an adapter is present for the object you plan to add. See the *vRealize Operations Manager vApp Deployment and Configuration Guide* .

Note   Objects added to vRealize Operations Manager via API will require an OSI license per object.

Procedure

1 In the menu, click **Administration**, then select **Configuration > Inventory** 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 as a host object type for the vSAN adapter type.

What to do next

When you add an individual object, vRealize Operations Manager does not begin collecting metrics for the object until you turn on data collection. See Inventory : List of Objects.

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. See Creating and Assigning Tags.

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.

Apart from the parent-child relationship, you can also define new relationships in vRealize Operations Manager . The relationship between objects in your environment can be one-to-many, many-to-one, or one-one, the relationship can be defined in horizontal , vertical, or diagonal levels.

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.

Object Relationships Workspace

Objects in an enterprise environment are related to other objects in that environment. Objects are either part of a larger object, or they contain smaller component objects, or both.

How Object Relationships Works

When you select a parent object, vRealize Operations Manager shows any related child objects. You can delete a child object or add more child objects from the list of objects in your environment.

Where You Find Object Relationships

At the Home page, select **Administration**. Then select **Configuration > Object Relationships** in the left pane.

Object Relationships Workspace Options

- Two columns in the center pane display the existing parent-child relationships. You use the object tag options above the left column to select a parent object.

- Two columns in the right pane list objects in your environment. You use the object tag options above the right column to select the object to add as a child.

Table 7‑1. Object Tag Options

Option	Description
Collapse all.	Closes all the tag group selections.
Deselect All.	Tags remain selected until deselected. Use this option to deselect all tags.

When a parent object has children, the parent selection shows the child objects and the child object options are active.

Table 7‑2. Child Object Options

Option	Description
Clear Selections.	Clear all child object selections.
Select All.	Select all child objects. To remove most child objects from the relationship, use this option then click the child objects you do not want to delete.
Remove Selected Children from Relationship.	Removes the selected children from the relationship.
Remove All Children from Relationship.	Select all children listed on the page and remove them from the relationship.
Per Page.	Number of children to list per page.
Search.	Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.

Use the list options to manage the objects to add as children.

Table 7‑3. List Options

Option	Description
Clear Selections.	Clear all object selections.
Select All.	Select all objects displayed.
Add All Objects to Parent.	Select all children listed on the page and add them to the parent.
Per page.	Number of objects to list per page.
Search.	Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.

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 7-4. 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.
Licensing	Tag values are License Groups found under **Home > Administration > Mangement > 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** 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 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** 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.

Manage Object Tags Workspace

A large enterprise can have thousands of objects. When objects are assigned to a tag, and you choose to display objects with that tag value, the objects are easier to find on the Inventory list.

Where You Find Manage Object Tags

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane.

Click the **Manage Tags** icon above the list of tags in the middle pane.

Manage Object Tags Options

The Manage Object Tags screen appears with previously created tags listed. In the left pane, you add tags. In the right pane, you add tag values.

- Click **Add a New Tag** and type a new tag name, or select a tag to delete.

- For the selected tag, click **Add a New Tag Value** and type a new tag value name, or select a tag value to delete.

- For the GEO Location tag, tag values are identified with a location on a world map. Select the tag value and click **Manage Location** to display the **Manage Location** map and pick a geographical location. Objects assigned to that tag value appear in that geographical location on the Inventory : Geographical Map of Objects.

Manage Object Type Tags Workspace

Every object in your environment is of a particular object type. You use Manage Object Type Tags to control the object type tags displayed.

How Manage Object Type Tags Works

For every adapter instance installed, vRealize Operations Manager discovers objects in your environment and starts collecting data from those objects.

Where You Find Manage Object Type Tags

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane. Click the **Manage Object Type Tags** icon above the list of tags.

Manage Object Type Tags Options

Depending on the number of adapters installed , there may be hundreds of object type tags. The Manage Object Type Tags options allow you to turn on or off the tags listed.

- Type a filter word to show the object type tags with the word.

- Name lists all the object type tags.

- To toggle the display of an object type tag, select the check box in the Show Tag column of its row.

Inventory : List of Objects

vRealize Operations Manager discovers objects in your environment for each adapter instance and lists them. From the complete list of all the objects in your environment, you can quickly access and configure any object. For example, you can check if a datastore is connected or providing data, or you can power on a virtual machine.

How the List Works

Objects appear in a data grid. To find a particular object, you can sort a column in the grid or search for a filter word. In addition to sorting and searching, assigning objects to object tags makes it easier to find objects and metrics.

Where You Find the List

In the menu, click **Administration**, then click **Inventory** . The system lists all the objects in your environment.

Inventory List Options

The center pane includes object tag options. The right pane includes toolbar options for all of the objects in your environment.

Table 7-5. Object Tag Options

Option	Description
Collapse all	Closes all the tag group selections.
Deselect All	Tags remain selected until deselected. Use this option to deselect all tags.

Table 7-5. Object Tag Options (continued)

Option	Description
Manage Tags	Add a tag or tag value. See Manage Object Tags Workspace.
Manage Object Type Tags	There might be many object type tags. Use this option to choose the object type tags to display. See Manage Object Type Tags Workspace.

Use the toolbar options to manage objects.

- Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.

- Select the object to manage from the list. If an object tag is selected, only objects of the selected tag value are listed. Column headings help you to identify the object. See Object List Widget.

Table 7-6. Inventory Toolbar Options

Option	Description
Action	Perform an action on the selected object. Available actions depend on the object type. For example, Power on VM applies to the selected virtual machine. See List of vRealize Operations Manager Actions
Open in external application	If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. For example, Open Virtual Machine in a vSphere Client or Search for VM logs in vRealize Log Insight.
Start Collecting	Turn on data collection for the selected object.
Stop Collecting	Do not collect data for the selected object. When data collection stops, vRealize Operations Manager retains metric data for the object in case data collection starts at a later time.
Perform Multi-Collecting	If an object collects metrics through more than one adapter instance, select the adapter instance or instances for data collection. Does not apply to objects that do not use the adapter instance.
Edit object	Edit the selected object. For example, add or change the maintenance schedule for a virtual machine. If multiple objects of the same type are selected, common identifiers for the object type are editable. For example, change the VM entity name of multiple datastores with a single edit. See Manage Objects Workspace.
Add object	vRealize Operations Manager discovers objects for most adapters. For adapters that do not support autodiscovery for all objects, the objects are manually added. See Manage Objects Workspace.

Table 7-6. Inventory Toolbar Options (continued)

Option	Description
Discover Objects	Perform an IP scan to discover objects associated with a particular adapter. See Discover Objects Workspace.
Delete object	Remove the object from the list.
Start maintenance	Take the object offline for maintenance. See Manage Maintenance Schedules for Your Object Workspace.
End maintenance	Terminate the maintenance period and put the selected object back online.
Clear Selections	Clear all object selections.
Select All	Select all objects displayed.
Show Detail	Display the **Summary** tab of the selected object.
Per page	The number of objects to list per page.

Manage Objects Workspace

To collect data from an object, you might need to add an object or edit an existing object in your environment. For example, you might need to add objects for an adapter that does not support autodiscovery, or change the maintenance schedule of an existing object.

Where You Find Manage Objects

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane. Click the plus sign to add an object or the edit icon to edit the selected object.

Items that appear in the window depend on the object that you are editing. Not all options can be changed.

Table 7-7. Manage Objects Add or Edit Options

Options	Description
Display name	Name of the object. Use only letters and numbers. Do not use nonalphanumeric characters or spaces.
Description	(Optional) For informational purposes only.
Adapter Type	If you are editing an object, you cannot change the adapter type.
Adapter Instance	If you are editing an object, you cannot change the adapter instance.
Object Type	If you are editing an object, you cannot change the object type. More configuration options might appear, depending on the object type.

Table 7-7. Manage Objects Add or Edit Options (continued)

Options	Description
Collection Interval	The collection interval for an object influences the collection status for the object. The collection interval for the adapter instance determines how often to collect data. For example, if the collection interval for an adapter instance is set to five minutes, setting the collection interval for an object to 30 minutes prevents the object from having the No Data Receiving collection status after five collection cycles or 25 minutes. In cases of adapter instances such as vRealizeOpsMgrAPI and HttpPost that push data to vRealize Operations Manager through the REST API, when data is no longer pushed, the status of the adapter instance is changed to Down after five collection intervals. For example, if the process pushes data every ten minutes and is stopped, the status of the adapter instance is changed to Down after 50 minutes. This behavior is expected for these adapter instance types.
Dynamic Thresholding	On by default, to enable dynamic thresholding and early warning smart alerts. See vRealize Operations Manager Dynamic Thresholds

Discover Objects Workspace

If vRealize Operations Manager does not discover objects after an adapter instance is configured, use manual discovery. Discovering objects is more efficient than adding objects individually.

Note You use discovery to define objects for embedded adapters. vRealize Operations Manager discovers objects that use external adapters.

Where You Find Discover Objects

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Discover Objects** in the List tool bar.

Discover Objects

The Discoveries section of the describe.xml file for the adapter might include parameters for discovery information. The describe.xml file is in the conf subfolder of the adapter, for example xyz_adapter3/conf/describe.xml.

Options	Description
Collector	Collector that vRealize Operations Manager uses to discover objects. Only the vRealize Operations Manager Collector is added during installation.
Adapter Type	Adapter type for the objects to discover.
Adapter Instance	Adapter instance of the selected adapter type.

Options	Description
Discovery Info	Selection depends on the adapter type. For example, for a vCenter adapter, the Discovery Info selection adds an option to discover objects of a particular object type.
Only New Objects	On by default, to omit objects that are already discovered.

Discovery Results List

When you use the Discover Objects feature to manually discover objects in your environment, vRealize Operations Manager lists the objects of the specified object type. You can choose the objects to monitor.

Where You Find Discovery Results

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Discover Objects** in the List tool bar.

After you make selections in the Discover Objects Workspace, click **OK**. With the default setting, vRealize Operations Manager displays only newly discovered objects. See .

Table 7-8. Object Types

Options	Description
Object Type	Discovered object types of the Object Type selected on the Discover Objects Workspace.
Object Count	Number of objects of the object type.
Import	When selected, imports the object type. Option is active and selectable for newly discovered object types.
Collect	When selected, imports the object type and starts collecting data. Option is active and selectable for newly discovered object types.
Credential	If the object type requires a login credential to collect data from the object., the value is **True**.

Double-click the Object Type to display a list of objects to monitor.

Table 7-9. Objects

Options	Description
Object	Objects of the selected type that exist in the environment for the adapter. For example, the vCenter adapter discovers objects in the vCenter Server system.
Import	When selected, imports the object but does not start collecting data. Option is active and selectable for newly discovered objects that do not exist in the vRealize Operations Manager environment .

Table 7-9. Objects (continued)

Options	Description
Exists	Indicates that the object exists in the vRealize Operations Manager environment.
Collect	When selected, imports the object and starts collecting data. Option is active and selectable for newly discovered objects that do not exist in the vRealize Operations Manager environment.

Manage Maintenance Schedules for Your Object Workspace

You use maintenance mode to take an object offline. Many objects in your environment might be intentionally taken offline. For example, you might deactivate a server to update software. If vRealize Operations Manager collects metrics when the object is offline, it might generate incorrect alerts that affect the data for the object's health. When an object is in maintenance mode, vRealize Operations Manager does not collect metrics from the object and does not generate alerts for it.

How Maintenance Schedules Work

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object into maintenance mode from midnight until 3 a.m. every Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can put an object in maintenance mode or take it out of maintenance mode, even if it has an assigned maintenance schedule.

Where You Find Manage Maintenance Schedules

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Start Maintenance** in the List tool bar.

Table 7-10. Manage Maintenance Schedules Options

Options	Description
I will come back and end maintenance myself.	Maintenance mode starts for the selected object when you click **OK**. You must manually end maintenance mode for this object.
End maintenance in	Type the number of minutes that the object is in maintenance mode.
End maintenance on	Click the calendar icon, and select the date that maintenance mode ends.

Define Custom Property Workspace

In vRealize Operations Manager , you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign custom properties to any subset of objects irrespective of the adapter

kind and resource kind. You can use a mouse click, search filter, or a tag selector to select the correct object.

Where You Find Add/Edit Custom Property

In the menu, select **Administration**, then click **Inventory** in the left pane. Click **Add/Edit Custom Property** in the List tool bar.

Table 7-11. Add/Edit Custom Property

Options	Description
Property Name	Select or enter a property name.
Type	Select the property type from the drop-down menu.
Value	Enter a value for the property.

You can assign the custom properties defined in this page to the Custom Object Groups and New Groups.

For more information, see Custom Object Groups Workspace to Create a New Group.

Inventory : Geographical Map of Objects

vRealize Operations Manager discovers objects in your environment for each adapter. Objects that are assigned a GEO Location tag appear on a geographical map. You can use this map to quickly locate your objects in the world.

How the Geographical Map Works

Objects with the GEO Location tag appear on a map of the world.

- To create a GEO Location tag, see Manage Object Tags Workspace.

- To assign objects to the tag, see Creating and Assigning Tags.

Where You Find the Geographical Map

In the menu, select **Administration**, then navigate to **Configuration > Inventory** in the left pane. Click the **Geographical** tab.

Geographical Map Options

Use the plus sign to zoom in. Use the minus sign to zoom out. Click and drag to pan the map to the left or right.

Managing Custom Object Groups in 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.

Object Group Types in vRealize Operations Manager

An object group type is an identifier that you apply to a specific group of objects in your environment to categorize them. You can add new group types, and apply them to groups of objects so that vRealize Operations Manager can collect data from the object group and display the results in the dashboards and views.

How the Group Types Work

Use group types to categorize your objects so that the system can apply policies to them to track, and display specific status, such as alerts, workload, faults, risk, and so on.

When you create a new group type, vRealize Operations Manager adds it to the existing list of group types, and creates a new folder with the name of your group type in the Environment Custom Groups list.

When you create a new group of objects, you assign a group type to that group of objects. You add objects from the inventory trees to your custom group, then create your dashboard, add widgets to the dashboard, and configure the widgets to display the data collected from the objects in the group. You can then monitor and manage the objects.

You can apply a group type to a group of objects that you create manually, or to object groups that you cannot modify, such those added by adapters. Each adapter that you add to vRealize Operations Manager adds one or more static groups of objects to group the data received from the adapter sources.

The list of group types appears in the Content area under Group Types. The custom object groups appear in the Environment area under Custom Groups.

Where You Create and Modify a Group Type

To create or modify a group type, click **Administration** in the menu, then **Configuration > Group Types** in the left pane.

Group Type Options

You can add, edit, or delete group types. You cannot edit group types that are created by adapters.

Groups Tab on the Environment Overview Pane

Groups are containers that can contain any number and type of objects in your environment. vRealize Operations Manager collects data from the objects in the group and displays the results in dashboards and views that you define.

How Groups Work

Groups are installed with vRealize Operations Manager , created by an adapter, or created by a user. Based on the group criteria, you can use groups to organize your environment and monitor all objects in the group together. You can also assign policies to groups and make group membership dynamic.

For example, if you have a set of vSphere hosts and you do not want to generate alerts when the host goes into maintenance mode, you can put the vSphere hosts in a group and assign a policy that includes a maintenance schedule setting. During the maintenance period, vRealize Operations Manager ignores any metrics for those objects and does not generate any alerts. After the maintenance period ends, vRealize Operations Manager returns to monitoring the objects and generates alerts if an outage occurs.

Where You Find Custom Groups

To access Custom Groups that you create, click **Environment** on the top menu, then click the **Custom Groups** tab.

Custom Group Options

Click the **ADD** button to add a group. You can only edit, clone, or delete a user-created group. You cannot modify groups installed with vRealize Operations Manager or by an adapter.

You can click the **Horizontal Ellipses** to import or export the custom group. The Groups data grid displays an overview of the state of each group. You can use the All Filters option to sort the custom groups based on Name, ID, Group Type, and Description columns.

To sort the list of custom groups based on columns, click the arrow next to the following columns:

- Name

- Health

- Risk

- Efficiency

- Description

- Members Count

Table 7-12. Group Data Grid Options

Option	Description
Name	Select the group name to display a summary of the group. Select to the right of the name to edit, clone, or delete the group.
Summary	Criticality of the health, risk, and efficiency of any group. Click a group with a red, orange, or yellow criticality to get more details about potential problems with objects in the group.
Members Count	Displays the number of members in the selected group.
Policy	Displays the policy associated with the selected group.
Dynamic Membership	Displays whether the group is static or dynamic. The available options are true and false.
Defined by	Displays who has defined the attributes of the group. The available options are: - System - User Defined - Management Pack

Custom Object Groups Workspace

You can create and edit custom groups of objects to have vRealize Operations Manager collect data from the objects and display the results in the dashboards and views so that you can monitor your objects and take action on them when problems occur.

How the Custom Groups Workspace Works

When you create a new object group, you define a meaningful group name, and select the group type. To associate the custom object group with a policy for analysis, you select the policy in the group creation wizard. You can leave the policy selection blank to not associate a policy with the object group. When the policy selection is blank, the custom object group is associated with the policy that is designated as the default policy.

You select the object types, and determine whether membership in the object group is static, dynamic, or a combination of static and dynamic membership.

- To create a static object group, you add objects to the group. You do not include criteria for object membership.

- To create a dynamic object group that vRealize Operations Manager updates based on specific criteria, you select the object type and define membership criteria for the group based on metrics, relationships, and properties.

When you add objects to a custom object group, a new folder appears in the Custom Groups navigation pane on the left, and includes the member objects.

Where You Create and Modify Object Groups

To create or modify static or dynamic object groups, or object groups that have a combination of static and dynamic membership, click **Environment > Custom Groups**. The **Custom Groups** tab displays a list of custom object groups, and the object groups for adapters added to vRealize Operations Manager .

To edit existing groups, select a group and click the edit icon on the **Custom Groups** tab.

Custom Object Groups Workspace to Create a New Group

You can create a new object group, define custom properties, assign a group type and objects to the group. When you create the group, you can assign a policy, or leave the policy selection blank to apply the default policy. vRealize Operations Manager collects data from the objects in the group based on the settings in the policy that is associated with the group. The results appear in the dashboards and views.

Where You Assign Custom Group Type, Policy, and Membership

To assign the group type, policy, and membership, click **Environment**, click **Custom Groups**, and click the plus sign to add a new group. In the New Group workspace, you can define the membership criteria, and select the objects to include or exclude.

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.

Table 7‑13. New Group Workspace

Option	Description
Name	Meaningful name of the object group.
Group Type	Categorization for the object group. New custom groups appear in a dedicated folder in the Custom Groups navigation pane on the left.
Policy	Assigns a policy to one or more groups of objects to have vRealize Operations Manager analyze the objects according to the settings in your policy, trigger alerts when the defined thresholds are violated, and display the results in dashboards, views, and reports. You can assign a policy to the group when you create the group, or you can assign it later from the edit custom group wizard or from the policies area.
Keep group membership up to date	For dynamic object groups, vRealize Operations Manager can discover objects that match the criteria for the group membership according to the rules that you define, and update the group members based on the search results.
Define Membership Criteria pane	Defines the criteria for a dynamic object group and has vRealize Operations Manager keep the object membership of the group current. ■ Object Type drop-down menu. Selects the type of objects to add to the group, such as virtual machines. ■ Metrics, Relationship, and Properties criteria drop-down menu. Defines the criteria for vRealize Operations Manager to apply to collect data from the selected objects. ■ Metrics. An instance of a data type, or attribute, that varies based on the object type. A metric is used as measurement criteria to collect data from objects. For example, you can select system attributes as a metric, where an attribute is a type of data that vRealize Operations Manager collects from objects. ■ Relationship. Indicates how the object is related to other objects. For example, you can require a virtual machine object to be a child object that contains a certain word in the vSphere Hosts and Clusters navigation tree. ■ Properties. Identifies a configuration parameter for the object. For example, you can require a virtual machine to have a memory limit that is greater than 100KB. ■ Add. Includes another metric, relationship, or property for the object type. ■ Remove. Deletes the selected object type from the membership criteria, or delete the selected metric, relationship, or property type from the criteria for the object type. ■ Reset. Resets the criteria for the first metric, relationship, or property that you define. ■ Adds another criteria set. Adds another object type to add to the group. For example, you might want to create a single object group to track vCenter Server instances and Host Systems. ■ Preview button. After you define the membership criteria, previews the list of objects in the group to verify that the criteria you defined is applicable to the group of objects. If the criteria that you defined is valid, the preview displays applicable objects. If the criteria is not valid, the preview does not display any objects.

Table 7-13. New Group Workspace (continued)

Option	Description
Objects To Always Include pane	Determine which objects to include in the group every time vRealize Operations Manager collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of vRealize Operations Manager , these objects were called a allowlist.

- Filtered objects pane. Displays the list of available object groups and the objects in each group. To always include objects in the group, select the check box for a group or select individual objects in a group, and click the **Add** button.
- **Add button**. Adds the selected objects to the right pane for permanent inclusion in the object group.
    - **Selected objects only**. Adds only the selected objects to the object group permanently.
    - **Selected objects and descendants**. Adds the selected object and the descendants of the selected objects to the object group permanently.
- Objects to always include (n) pane. Lists the objects that you add to the include list. You must select the check box in the right pane to confirm inclusion of the objects. The number of objects selected for inclusion is reflected by the (n) variable in the title of the pane.
- Remove button. Removes the objects selected in the right pane from the list of objects to always include.
    - **Selected objects only**. Removes only the selected objects from the list of objects to always include.
    - **Selected objects and direct children**. Removes the selected objects and the children of the selected objects from the list of objects to always include.
    - **Selected objects and all descendants**. Removes the selected objects and the descendants of the selected objects from the list of objects to always include.

Table 7-13. New Group Workspace (continued)

Option	Description
Objects To Always Exclude pane	Determine which objects to exclude from the group every time vRealize Operations Manager collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of vRealize Operations Manager , these objects were called a denylist. ■ Filtered objects pane. Displays the list of available object groups and the objects in each group. To always exclude objects from the group, select the check box for a group or select individual objects in a group, and click the **Add** button. ■ **Add button**. Adds the selected objects to the right pane for permanent exclusion from the object group. 　■ **Selected objects only**. Adds only the selected objects to be permanently excluded from the object group. 　■ **Selected objects and descendants**. Adds the selected objects and the descendants of the selected objects for permanent exclusion from the object group. ■ Objects to always exclude (n) pane. Lists the objects that you add to the exclude list. You must select the check box in the right pane to confirm exclusion of the objects. The number of objects selected for exclusion is reflected by the (n) variable in the title of the pane. ■ Remove button. Removes the objects selected in the right pane from the list of objects to always exclude. 　■ **Selected objects only**. Removes only the selected objects from the list of objects to always exclude. 　■ **Selected objects and direct children**. Removes the selected objects and the children of the selected objects from the list of objects to always exclude. 　■ **Selected objects and all descendants**. Removes the selected object and the descendants of the selected objects from the list of objects to always exclude.
Assign Custom Properties	In vRealize Operations Manager , you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign the newly defined custom properties to new groups or existing groups. ■ Property Name. Select or specify a name for the custom property. ■ Type. Select the type of custom property from the drop-down menu. 　The custom property can either be a string or a numeric. ■ Inclusion Value. Specify a custom property value, which should be assigned to this custom property when an object is added to the group. ■ Exclusion Value. Specify a custom property value, which should be assigned to this custom property when an object leaves the group. ■ **Reset**. Resets the custom property to a non-zero value. ■ **Remove**. Removes the custom property from the group. ■ **Add Another Custom Property**. Adds another custom property to the group.

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.

Applications Tab on the Environment Overview Pane

Applications are groups of related objects in your environment that mimic an application in your business. Use the summary to track the health of objects in the application and help troubleshoot performance issues.

How Applications Work

In vRealize Operations Manager , each application contains one or more tiers and each tier contains one or more objects. The tier is a convenient way to organize objects that perform a specific task in an application. For example, you can group all of your database servers together in a tier.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the application.

Construct an application to view a particular segment of your business. The application shows how the performance of one object affects other objects in the same application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data for your business, you see a yellow, orange, or red status if the application health is degrading. Starting with the application summary dashboard, you can investigate which server is causing or exhibiting the problem.

Where You Find Applications

In the menu, click **Environment**, then click the **Applications** tab.

Applications defined in a previous release of vRealize Operations Manager appear after an upgrade.

Application Options

Select an application to edit or delete, or click the **ADD** button to add an application.

The Applications data grid displays an overview of the state of each application.

Table 7-14. Application Data Grid Options

Option	Description
Name	Select the application name to display a summary of the application. Select to the right of the name to edit or delete the application.
Summary	Criticality of the health, risk, and efficiency of any application. Click an application with a red, orange, or yellow criticality to see more details about potential problems with objects in the application.

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 **ADD** button.

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 see Evaluate Object Information Using Badge Alerts and Summary Tab .

To investigate the source of the problem, click the application name and evaluate the object summary information. See the *vRealize Operations Manager User Guide* .

Add Application

When you add an application to an environment, you select from a list of predefined templates or create your own custom template, to group the objects to monitor in your application.

Where You Find Add Application

In the menu, click **Environment**, then **Groups and Applications > Applications** in the left pane. On the **Applications** tab, click the plus sign.

Add Applications Options

Each predefined template provides you with a list of suggested tiers designed to help you group related objects that perform a specific task in your application. After you select an option, you can alter the selection and number of tiers on the Application Management page.

Option	Description
Basic n-tier Web App	Use this template for any basic application.
Advanced n-tier Web App	Use this template for an application that monitors more physical devices, such as the devices that vRealize Operations Manager discovers when you add a network-related Management Pack or Management Packs.
Legacy non-Web App	Use this template for an application that has no Web-related objects.
Network	Use this template for an application that has only network-related objects.
Custom	Select this option to build your own application topology.

Application Management Dialog Box

You use Application Management to select the objects for your application. The objects you select are grouped in tiers and help you to track the health of your application.

Where You Find Application Management

In the menu, click **Environment**, then click the **Groups and Applications** menu and select **Applications**. On the **Applications** tab, click the plus sign. After you select an application template, click OK.

Application Management Options

At the top of the screen, enter a new application name or use the default name from the Add Application page. The application name must be unique.

Below the name, the page is divided into the tier row and the objects row. On each row, selections in the pane on the left filter the selections in the pane on the right.

The tier row is where you select the tiers to populate with objects to monitor for the application.

Table 7-15. Tier Row

Option	Description
Tiers pane	Select the tier where you want to place your objects. You can add or delete tiers to fit your application.
Tier Objects pane	Add or remove objects that serve a common function and to monitor. For example, to monitor all the virtual machines that are database servers for the application, put them in the database tier.

The object row is where you select objects to add to the tiers.

Table 7-16. Object Row

Option	Description
Object Tags pane	Expand a tag to see a group of objects with that tag value. For example, if Adapter Types is an object tag, the tag values include vCenter Adapter, and an object is an adapter instance. Objects are not displayed. The tag filters the object pane. To select a tag value, click once. To deselect a tag value, click twice. Tag values remain selected until they are deselected.
Objects pane	Drag an object with the object tag value to add to the Tier Objects pane. To find an object, search by name. Each object listed includes identifier information to help distinguish between objects of similar names. **Add All Objects To Parent** adds all the objects to a tier.

Configuring Data Display 8

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 8-1. 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/opscli` 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	`/usr/lib/vmware-vcops/tools/opscli`

3 Run the import command.

Operating System	File Location
vApp	`./ops-cli.sh file import reskndmetric YourCustomFilename.xml`

   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 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 8-2. 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 configure, 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.

Table 8-2. Summary of Widgets (continued)

Widget Name	Description
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.
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 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 resource 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.

Table 8-2. Summary of Widgets (continued)

Widget Name	Description
Workload	Shows workload information for a selected resource.
Workload Pattern	Shows a historical view of the hourly workload pattern of an object.
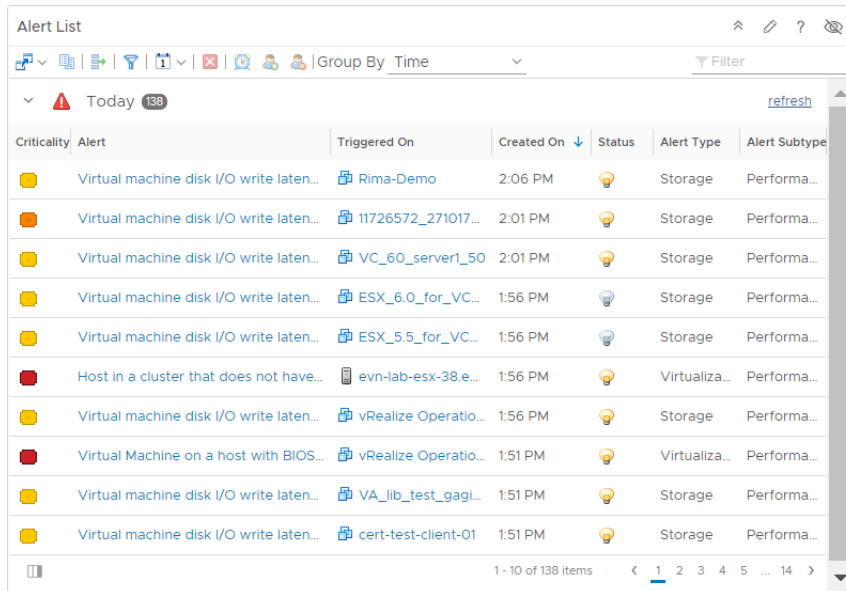
For more information about the widgets, see the vRealize Operations Manager help.

Alert List Widget

The Alert List widget is a list of alerts for the objects it is configured to monitor. You can create one or more alert lists in vRealize Operations Manager for objects that you add to your custom dashboards. The widget provides you with a customized list of alerts on objects in your environment.

How the Alert List Widget and Configuration Options Work

You can add the Alert List widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. You edit an Alert List widget after you add it to a dashboard. The changes you make to the options create a custom alert list to meet the needs of the dashboard users.

Where You Find the Alert List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Alert List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	Actions you can run on the selected alert.
	For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems.
Reset Interaction	Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.
	Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards.
Perform Multi-Select Interaction	If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.
	Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to enable the interaction.
Display Filtering Criteria	Displays the object information on which this widget is based.
Select Date Range	Limits the alerts that appear in the list to the selected date range.
Cancel Alert	Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.
	You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again.

Option	Description
Suspend	Suspend an alert for a specified number of minutes.
	You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.
	The user who suspends the alert becomes the assigned owner.
Take Ownership	As the current user, you make yourself the owner of the alert.
	You can only take ownership of an alert, you cannot assign ownership.
Release Ownership	Alert is released from all ownership.
Group By	Group alerts by the options in the drop-down menu.
Filter	Locate data in the widget.

Table 8-3. Group By Options

Option	Description
None	Alerts are not sorted into specific groupings.
Time	Group alerts by time triggered. The default.
Criticality	Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the Alert List Widget Data Grid table.
Definition	Group alerts by definition, that is, group like alerts together.
Object Type	Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together.

Alert List Widget Data Grid Options

The data grid provides information on which you can sort and search.

Expand the grouped alerts to view the data grid.

Option	Description
Criticality	Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon.
	The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.
Alert	Description of the alert.

Option	Description
Triggered On	Name of the object for which the alert was generated.
Created On	Date and time when the alert was generated.
Status	Current state of the alert.
Alert Type	Alert type is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution. The possible values include: ■ Application ■ Virtualization/Hypervisor ■ Hardware (OSI) ■ Storage ■ Network
Alert Sub-Type	Alert subtype is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution. The possible values include: ■ Availability ■ Performance ■ Capacity ■ Compliance ■ Configuration
Importance	Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm.

Alert List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Objects	Select objects on which you want to base the widget data.
	1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	

Option	Description
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
	If the objects have an input transformation applied, you select tag values for the transformed objects.

Option	Description
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.
Alert Related	A group of filters limits the alerts that appear in this alert list to those that meet the selected criteria.
	If the objects on which the alerts are based have an input transformation applied, you define filters for the alerts based on the transformed objects.
	You can configure the following filters:
	■ Alert Type. Select the subtype in the type list. This value was assigned when you configured the alert definition.
	■ Status. Select one or more alert states to include in the list.
	■ Control State. Select one or more control states to include in the list.
	■ Criticality. Select one or more levels of criticality.
	■ Impact. Select one or more alert badges to include in the list.

Alert Volume Widget

The Alert Volume widget is a trend report for the last seven days of alerts generated for the objects it is configured to monitor in vRealize Operations Manager. You can create one or more alert volume widgets for objects that you add to your dashboards. The alert volume provides you with a customized trend report on objects that helps you identify changes in alert volume, indicating a problem in your environment.

How the Alert Volume Widget and Configuration Options Work

You can add the Alert Volume widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. The changes you make to the options create a custom widget to meet the needs of the dashboard users.
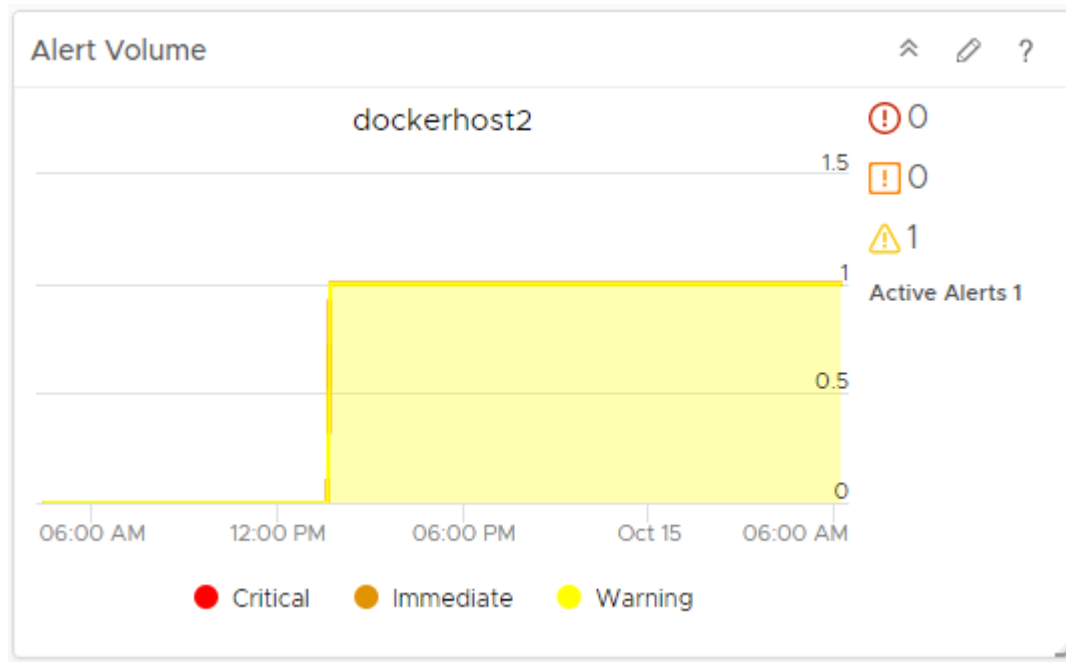
Where You Find the Alert Volume Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Alert Volume Widget Display Options

The Alert Volume widget displays a trend chart, symptoms by criticality, and active alerts.

Option	Description
Trend chart	Volume of critical, immediate, and warning symptoms for the configured objects.
Symptoms by criticality	Number of symptoms for each criticality level.
Active Alerts	Number of active alerts. Alerts can have more than one triggering symptom.

Alert Volume Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Anomalies Widget

The Anomalies widget displays the anomalies for a resource for the past 6 hours at time intervals you set.

The Anomalies widget shows or hides time periods when the metric violates a threshold that configured. The widget color indicates the criticality of the violation.
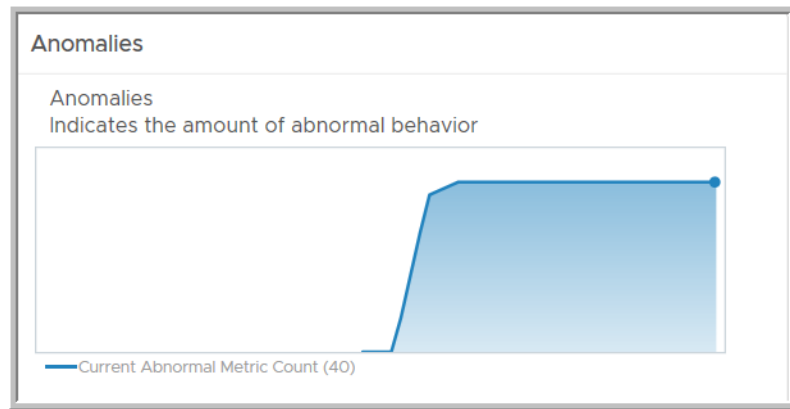
Where You Find the Anomalies Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Anomalies Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Anomaly Breakdown Widget

The Anomaly Breakdown widget shows the likely root causes for symptoms for a selected resource.

How the Anomaly Breakdown Widget and Configuration Options Work
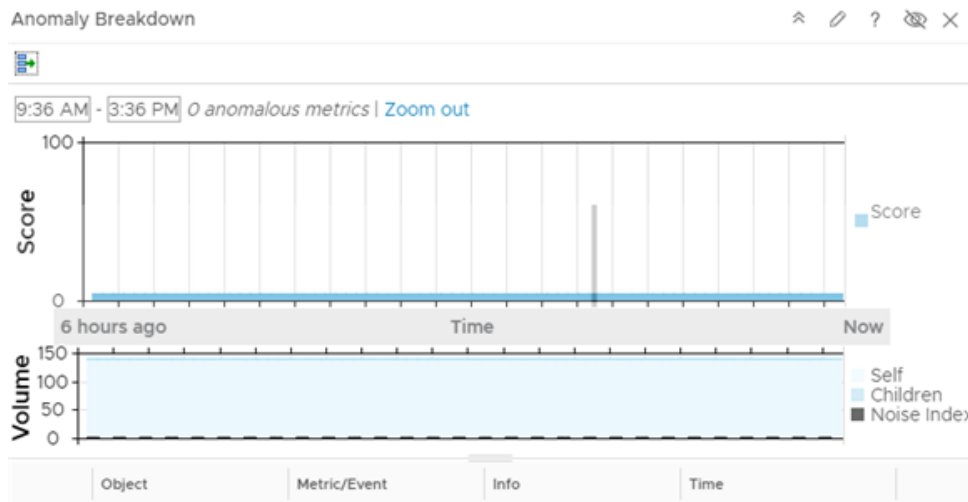
You can add the Anomaly Breakdown widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Where You Find the Anomaly Breakdown Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Anomaly Breakdown Widget Display Options

The Anomaly Breakdown widget displays scores, volume, and a list of anomaly metrics.

Option	Description
Score	Anomaly value.
Volume	vRealize Operations Manager full set metric count for the selected object in the specified time range.
Anomaly Metrics List	List of alarms for the selected object in the specified time range.

Anomaly Breakdown Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Show Bar Details	If the widget is displaying data for multiple objects, you can select a row and click this button to view the list of alarms for the selected object.
Perform Multiple Interaction	If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.
	Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to enable the interaction.

Anomaly Breakdown Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Mode	Display a single object or multiple objects.
Show	Select the number of objects to display in multiple objects mode.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.

Container Details Widget

The Container Details widget displays graphs that show a summary of child objects, metrics, and alerts of an object in the inventory.
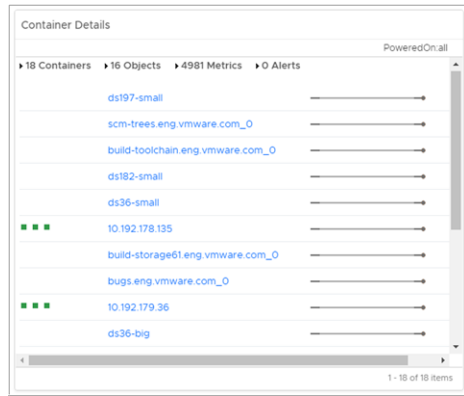
How the Container Details Widget and Configuration Options Work

The Container Details widget treats objects from the inventory as containers and objects. Containers are objects that contain other objects. The widget lists the containers and shows the number of containers, objects, metrics, and alerts of the observed object. The widget also displays the alerts of each container and an icon links to its child objects. For example, if you select from the inventory a host that contains three objects such as, two virtual machines and one datastore, the Container Details widget displays summary information with three containers, two objects that are the child objects of the two virtual machines, and the number of alerts for the host and the number of metrics for the child objects of the host. The widget also lists each of the three containers, with the number of alerts for each object. Clicking an object in the graph takes you to the object details page. When you point to the icon next to the object, a tool tip shows the name of the related resource and its health. For example, when you point to the icon next to a virtual machine, the tool tip shows a related datastore and its health. Clicking the icon takes you to the object detail page of the related object, which is the datastore following the example.

You edit a container details widget after you add it to a dashboard. You can configure the widget to take information from another widget in the dashboard and to analyze it. When you select **Off** from the Self Provider option and set source and receiver widgets in the **Widget Interactions** menu during editing of the dashboard, the receiver widget shows information about an object that you select from the source widget. For example, you can configure the Container Details widget to display information about an object that you select from the Object Relationship widget in the same dashboard.

Where You Find the Container Details Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Container Details Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Mode	You can change the size of the graph using the Compact or Large buttons.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Capacity Remaining Widget

The Capacity Remaining widget displays a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.

Where You Find the Capacity Remaining Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Capacity Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Container Overview Widget

The Container Overview widget gives a graphical presentation of the health, risk, and efficiency of an object or list of objects in the environment.

How the Container Overview Widget and Configuration Options Work

The Container Overview widget displays the current status, the status for a previous time period of the health, risk, and the efficiency of an object or list of objects. You can configure the widget to display information for one or more objects that you are interested in when you select the **Object** mode during configuration of the widget. The widget displays information for all objects from an object type or types when you select the **Object Type** mode during configuration of the widget. You can open the object detailed page of each object in the data grid when you click the object.

You edit a container overview widget after you add it to a dashboard. You can configure the widget to display information about an object or to display information about all objects from an object type by using the **Object** or **Object Type** mode. The configuration options change depending on your selection of mode.

Where You Find the Container Overview Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Container Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about other widgets or dashboards.

Option	Description
Perform Multi-Select Interaction	If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items. Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to enable the interaction.
Filter	You can filter the objects in the data grid.
Dashboard Navigation	You can explore information from another dashboard. **Note** This toolbar icon exists when you configure the widget to interact with a widget from another dashboard. Use **Dashboard Navigation** menu during dashboard configuration to configure the widgets to interact. When you select an object from an object data grid and click the toolbar icon, it takes you to a related dashboard. For example, you can configure the widget to send information to a Topology Graph widget that is on another dashboard, for example dashboard 1. When you select a VM from the data grid, click **Perform Multi-Select Interaction** , click **Dashboard Navigation** and select **Navigate > dashboard 1**. It takes you to dashboard 1, where you can observe selected VM and objects related to it.

Container Overview Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Name	Name of the object
Health	Shows information about the health parameter. Status displays the badge of the current health status of an object. You can check the status in a tool tip when you point to the badge. Last 24 Hours displays the statistic of health parameter for last 24 hours.

Option	Description
Risk	Shows information about the risk parameter.
	Status displays the badge of the current risk status of an object. You can check the status in a tool tip when you point to the badge.
	Last Week displays the statistics of the health parameter for the last week.
Efficiency	Shows information about the efficiency parameter.
	Status displays the badge of the current efficiency status of an object. You can check the status in a tool tip when you point to the badge.
	Last Week displays statistic of the efficiency parameter for the last week.

Container Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Mode	Use **Object** to select an object from the environment to observe.
	Use **Object Type** to select the type of the objects to observe.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
Object Type	Select an object type in your environment on which you want to base the widget data. 1 Click the **Add Object Type** icon to search for and add an object type. When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. 2 Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type.

Current Policy Widget

The Current Policy widget displays the active operational policy that is assigned to your object or object group. vRealize Operations Manager uses the assigned policy to analyze your objects, control the data that is collected from those objects, generate alerts when problems occur, and display the results in the dashboards.

How the Current Policy Widget and Configuration Options Work

You add the Current Policy widget to a dashboard so that you can quickly see which operational policy is applied to an object or object group. To add the widget to a dashboard, you must have access permissions associated with the roles assigned to your user account.

The configuration changes that you make to the widget creates a custom instance of the widget that you use in your dashboard to identify the current policy assigned to an object or object group. When you select an object on the dashboard, the policy applied to the object appears in the Current Policy widget, with an embedded link to the policy details. To display the inherited and local settings for the applied policy, click the link.

Where You Find the Current Policy Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Current Policy Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. For example, to view the policy applied to each object that you select in the Object List widget, select **Off** for Self Provider.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Data Collection Results Widget

The Data Collection Result widget shows a list of all supported actions specific for a selected object. The widget retrieves data specific to a selected object actions and uses the action framework to run data collection actions.

How the Data Collection Results Widget and Configuration Options Work

You can add the Data Collection Results widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The Data Collection Results widget is a receiver of a resource or metric ID. It can interact with any resource or metric ID that provides widgets such as Object List and Metric Picker. To use the widget, you must have an environment that contains the following items.

■ A vCenter Adapter instance

■ A vRealize Operations Manager for Horizon View Adapter

■ A vRealize Operations Manager for Horizon View Connection Server

You edit a Data Collection Result widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where You Find the Data Collection Results Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Data Collection Results Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Results	Shows all finished and currently running actions for the selected object.
Choose Action	Shows a list with all supported actions specific for the selected object. The selected object is a result of widget interactions.

Data Collection Results Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget updates only when you open the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Config	Specifies self provider choice and selection of a resource instance.
Selected Object	When you select an object, this text box is populated by the object.

Option	Description
Start new data collection on interaction change	Indicates whether to start a new data collection action when the object selection changes in the source widget.
Objects	List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget.
Defaults	Specifies the default data collection action selected for each object type.
Object Types	List of object types in your environment that you can search or sort by column so that you can locate the object type on which you are basing the data that appears in the widget. You can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.
Default Data Collection Action	This panel is populated by the object type that you select in the object types list. You can select only one default data collection action for an object type.

DRS Cluster Settings Widget

The DRS Cluster Settings widget displays the workload of the available clusters and the associated hosts. You can change the Distributed Resource Scheduler (DRS) automation rules for each cluster.

How the DRS Cluster Settings Widget and Configuration Options Work

You can view CPU workload and memory workload percentages for each of the clusters. You can view CPU workload and memory workload percentages for each host in the cluster by selecting a cluster in the data grid. The details are displayed in the data grid below. You can set the level of DRS automation and the migration threshold by selecting a cluster and clicking **Cluster Actions > Set DRS Automation**.

You edit a DRS Cluster Settings widget after you add it to a dashboard. To configure the widget, click the edit icon at the upper-right corner of the widget window. You can add the DRS Cluster Settings widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The DRS Cluster Settings widget appears on the dashboard named vSphere DRS Cluster Settings, which is provided with vRealize Operations Manager.

Where You Find the DRS Cluster Settings Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

DRS Cluster Settings Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Cluster Actions	Limits the list to actions that match the cluster you select.
Show	The drop-down menu displays the parent vCenter Server instances where the clusters reside. You can also view the data centers under each parent vCenter Server instance. Select a parent vCenter Server to view the workload of the available clusters in the data grid. The default setting displays the clusters across all vCenters.
Filter	Filters the data grid by name, data center, vCenter, DRS settings, and migration threshold.

DRS Cluster Settings Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Name	Displays the names of the clusters in the selected parent vCenter Server instance.
Datacenter	Displays the data centers that belong to each cluster.
vCenter	Displays the parent vCenter Server instance where the cluster resides.
DRS Settings	Displays the level of DRS automation for the cluster. To change the level of DRS automation for the cluster, select **Cluster Actions > Set DRS Automation** from the toolbar. You can change the automation level by selecting an option from the drop-down menu in the Automation Level column.
Migration Threshold	Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster.
CPU Workload %	Displays the percentage of CPU in GHz available on the cluster.
Memory Workload %	Displays the percentage of memory in GB available on the cluster.

DRS Cluster Settings Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Efficiency Widget

The efficiency widget is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in vRealize Operations Manager usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.

How the Efficiency Widget and Configuration Options Work

You can add the efficiency widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning efficiency alerts generated over time, if the monitored object is a group.

- A trend line displays the efficiency status of the monitored object over time if the object does not provide its resources to any other object, or where no other object depends on the monitored object's resources. For example, if the monitored object is a virtual machine or a distributed switch.

- A pie chart displays the reclaimable, stress, and optimal percentages for the virtual machines that are descendants of the monitored object for all other object types. You use the chart to identify objects in your environment from which you can reclaim resources. For example, if the object is a host or datastore.

If the **Badge Mode** is set to **On**, only the badge appears.

Edit an efficiency widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Efficiency Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Efficiency Widget Display Options

The Efficiency widget displays an efficiency badge. The widget also displays an efficiency trend when not in badge mode.

Option	Description
Efficiency Badge	Status of the objects configured for this instance of the widget. Click the badge to open the **Alerts** tab for the object that provides data to the widget.
Efficiency Trend	Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears.

Efficiency Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Badge Mode	Determines whether the widget displays only the badge, or the badge and a weather map or trend chart. Select one of the following options: ■ On. Only the badge appears in the widget. ■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Environment Widget

The Environment widget displays the resources for which collects data. You can create one or more lists in vRealize Operations Manager for the resources that you add to your custom dashboards.

How the Environment Widget and Configuration Options Work

The Environment widget lists the number of resources by object or groups them by object type. You can add the Environment widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an Environment widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.
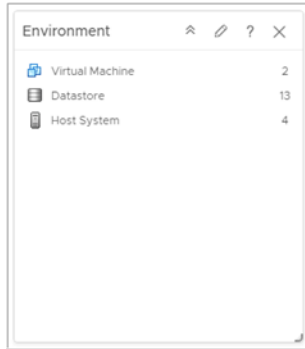
Where You Find the Environment Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Environment Overview Widget

The Environment Overview widget displays the health, risk, and efficiency of resources for a given object from the managed inventory.
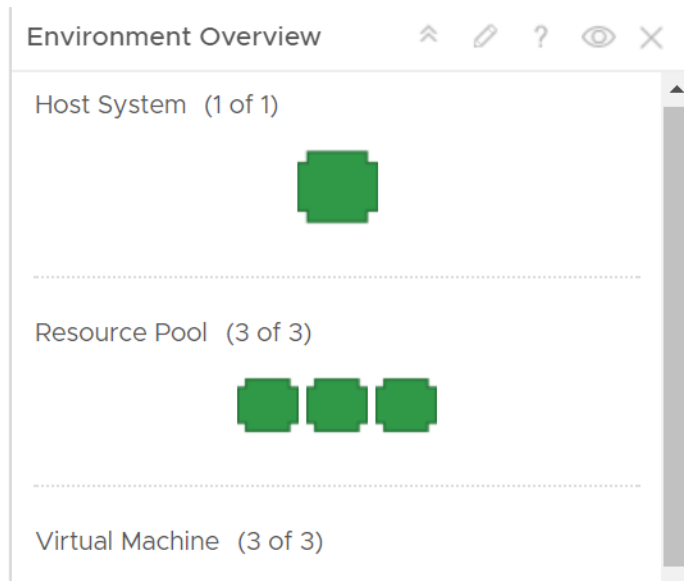
How the Environment Overview Widget and Configuration Options Work

You can add the Environment Overview widget to one or more custom dashboards.

The widget displays data for objects from one or several types. The data that the widget displays depends on the object type and category that you selected when you configured the widget.

The objects in the widget are ordered by object type.

The parameters for the health, risk, and efficiency of an object appear in a tool tip when you point to the object.

When you double-click an object on the Environment Overview widget, you can view detailed information for the object.

To use the Environment Overview widget, you must add it to the dashboard and configure the data that appears in the widget. You must select at least one badge and an object. Additionally, you can select an object type.

The Environment Overview widget has basic and advanced configuration options. The basic configuration options are enabled by default.

To use all features of the Environment Overview widget, you must change the default configuration of the widget. Log in to the vRealize Operations Manager machine and set `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. The change is propagated after you use the `service vmware-vcops-web restart` command to restart the UI.

You must use the **Badge** tab to select the badge parameters that the widget shows for each object. You must use the **Config** tab to select an object or object type. To observe a concrete object from the inventory, you can use the **Basic** option. To observe a group of objects or objects from different types, you must use the **Advanced** option.

Where You Find the Environment Overview Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about badges.

Option	Description
Badge	You can select a Health, Risk, or Efficiency badge for objects that appear in the widget. The tool tip of a badge shows the standard name of the badge.
Status	You can filter objects based on their badge status and their state.
Sort	You can sort objects by letter or by number.

Environment Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Selected Object	Object that is the basis for the widget data. To populate the text box, select **Config > Basic** and select an object from the list.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Badge	Defines a parameter to observe. You can select or deselect Health, Risk, and Efficiency parameters using check boxes. Default configuration of the widget selects all badges. Select at least one badge parameter.

Option	Description
Config	**Basic** List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget.
	Advanced You can use Object Types to select a type of the objects to observe information about health, risk, and efficiency. Double-click the object type to select it. Use the **Adapter Type** drop-down menu to filter the objects types based on an adapter. You can use the **Use vSphere Default** button to observe the main vSphere object types. To remove an object type from the list, click **Remove Selected** next to **Use vSphere Default**. You can use the **Object Type Categories** menu to select a group or groups of object types to observe. You can use the Object tree to select an object to filter the displayed objects. For example, to observe a datastore of a VM, double-click **Datastore** from the **Object Types** menu to select it. Click the datastore when it is in the list of object types, and find the VM in the object tree and select it. To return to your previous configuration of the widget, click **Datastore** from the list of object types and click **Deselect All** in the object tree window. The metrics tree and badge data grids are available configuration options only if the default configuration of the widget is changed. To use these configuration options, log in to the vRealize Operations Manager machine and set `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder.

Environment Status Widget

The Environment Status widget displays the statistics for the overall monitored environment.
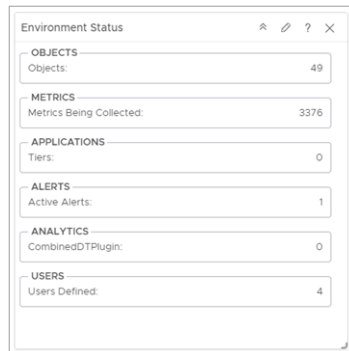
How the Environment Status Widget and Configuration Options Work

You customize the output of the widget by choosing a category such as Objects, Metrics, Applications, Alerts, Analytics, and Users. You can filter the data by using the tags tree from **Select which tags to filter** in the configuration window.

You edit an environment status widget after you add it to a dashboard. To configure the widget, click the pencil at the right corner of the widget window. You must select at least one type of information from **OBJECTS**, **METRICS**, **APPLICATIONS**, **ALERTS**, **ANALYTICS**, **USERS** categories for the widget to display. By default, the widget displays statistics information about all objects in the inventory. You can use the Select which tags to filter option to filter the information. The widget can interact with other widgets in the dashboard, taking data from them and displaying statistics . For example, you can have a Object List widget , which is the source of the data and an Environment Status widget, which is the destination. If you select objects and perform a multiselection interaction from the Object List widget, the Environment Status widget results are updated based on the selections you made in the Object List.

Where You Find the Environment Status Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Status Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
	The widget is also updated when it is in interaction mode. For example, when an item is selected in the provider widget, the content of the Environment Status widgets is refreshed.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Objects	Select objects on which you want to base the widget data.
	1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.

Option	Description
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Objects	The widget shows summarized information about the objects in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of resources. For example, if you select **Adapter Types > Container** from Select which tag to filter and click **Objects** and **Objects Collecting** , the widget displays the number of containers and collecting containers.
Metrics	The widget shows summarized information about available metrics. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of metrics.
Applications	The widget shows summarized information about available applications. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of applications.
Alerts	The widget shows summarized information about alerts in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of alerts.
Analytics	The widget shows summarized information about the analytics plug-ins. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of analytics.
Users	The widget shows the number of users defined in vRealize Operations Manager. Select **Administration > Access Control > User Accounts**.
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
	If the objects have an input transformation applied, you select tag values for the transformed objects.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Faults Widget

The Faults widget displays detailed information about faults experienced by an object

The Faults widget configuration options are used to customize each instance of the widget that you add to your dashboards.

Where You Find the Faults Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Faults Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Forensics Widget

The Forensics widget shows how often a metric has a particular value as a percentage of all values, within a given time period. It can also compare percentages for two time periods.

How the Forensics Widget and Configuration Options Work

You can add the Forensics widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit the Forensics widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where you Find the Forensics Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Forensics Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Percentile	Indicates how much data is above or below the specific value. For example, it indicates that 90% of the data is more than 4 when a vertical line occurs on the value 4.

Option	Description
Input Data	
	Select metrics on which you want to base the widget data. You can select an object and pick its metrics.
	1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.
	The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.
	While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.
	Click the **Select All** icon to select all the metrics in the list.
	Click the **Clear Selection** icon to clear your selection of metrics in the list.

Geo Widget

If your configuration assigns values to the Geo Location object tag, the geo widget shows where your objects are located on a world map. The geo widget is similar to the **Geographical** tab on the Inventory page.

How the Geo Widget and Configuration Options Work

You can move the map and zoom in or out by using the controls on the map. The icons at each location show the health of each object that has the Geo Location tag value. You can add the geo widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a Geo widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

Where You Find the Geo Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Geo Widget Toolbar Options

Option	Description
Zoom in	Zooms in on the map.
Zoom out	Zooms out on the map.

Geo Widget Configuration Options

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Heatmap Widget

The Heatmap widget contains graphical indicators that display the current value of two selected attributes of objects of tag values that you select. In most cases, you can select only from internally generated attributes that describe the general operation of the objects, such as health or the active anomaly count. When you select a single object, you can select any metric for that object.

How the Heatmap Widget and Configuration Options Work

You can add the Heatmap widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The Heatmap widget has a General mode and an Instance mode. The General mode shows a colored rectangle for each selected resource. In the Instance mode, each rectangle represents a single instance of the selected metric for an object.

You can click a color or the size metric box in the bottom of the Heatmap widget to filter the display of cells in the widget. You can click and drag the color filter to select a range of colors. The Heatmap widget displays cells that match the range of colors.

When you point to a rectangle for an object, the widget shows the resource name, group-by values, the current values of the two tracked attributes, virtual machine details, the metric name, and the value of the color. Click **Show Sparkline** to view the value.

You edit a Heatmap widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Heatmap Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Heatmap Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	Actions you can run on the selected alert.
	For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems.
Group Zoom	You can roll-up non-significant resources with similar characteristics into groups to obtain only the relevant data among the thousands of resources in the system. The roll-up method improves performance and decreases the memory usage. The roll-up box encompasses the average color and the sum of the sizes of all the resources. You can view all the resources by zooming in the roll-up box.
Show/Hide Text	Show or hide the cell name on the heatmap rectangle.
Show Details	If you configure the Heatmap widget as a provider to another widget, such as the Metric Chart widget , you can double-click a rectangle to select that object for the widget. If the widget is in Metric mode, double-clicking a rectangle selects the resource associated with the metric and provides that resource to the receiving widget. Optionally, you can select a cell from the heatmap and click the **Show Details** icon to see details about the cell.

Option	Description
Reset Interaction	Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.
Reset Zoom	Resets the heatmap display to fit in the available space.
Heatmap Configuration Drop-down	Select from a list of predefined heatmaps.

Heatmap Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	

Realize Operations Manager Configuration Guide

Option	Description
Objects	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Configurations	List of saved heatmap configuration options. You can create a configuration and save it in the list. From the options on the right, you can also delete, clone, and reorder the configurations.
Name	Name of the widget.
Group by	First-level grouping of the objects in the heatmap.
Then by	Second-level grouping of the objects in the heatmap.
Relational Grouping	After you select the Group by and Then by objects, select the **Relational Grouping** check box to reorganize the grouping of the objects, and to relate the objects selected in the Group by text box with the objects selected in the Then by text box.

Option	Description
Mode	**General mode** The widget shows a colored rectangle for each selected resource. The size of the rectangle indicates the value of one selected attribute. The color of the rectangle indicates the value of another selected attribute. **Instance mode** Each rectangle represents a single instance of the selected metric for a resource. A resource can have multiple instances of the same metric. The rectangles are all the same size. The color of the rectangles varies based on the instance value. You can use instance mode only if you select a single resource kind.
Object Type	Object that is the basis for the widget data.
Size by	An attribute to set the size of the rectangle for each resource. Resources that have higher values for the Size By attribute have larger areas of the widget display. You can also select fixed-size rectangles. In most cases, the attribute lists include only metrics that vRealize Operations Manager generates. If you select a resource kind, the list shows all the attributes that are defined for the resource kind.
Color by	An attribute to set the color of the rectangle for each resource.
Solid Coloring	Select this option to use solid colors instead of a color gradient. By default, the widget assigns red color for high value, brown color for intermediate value and green color for low value. Click the color box to set a different color for the values. You can add up to seven color thresholds by clicking color range.
Color	Shows the color range for high, intermediate and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes. By default, green indicates a low value and red indicates the high end of the value range. You can change the high and low values to any color and set the color to use for the midpoint of the range. You can also set the values to use for either end of the color range, or let vRealize Operations Manager define the colors based on the range of values for the attribute. If you leave the text boxes blank, vRealize Operations Manager maps the highest and lowest values for the **Color By** metric to the end colors. If you set a minimum or maximum value, any metric at or beyond that value appears in the end color.

Option	Description
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. If the objects have an input transformation applied, you select tag values for the transformed objects.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types. If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied. If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects. 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click **Add**. 5 To add another filter criteria set, click **Add another criteria set**.

Health Widget

The Health widget is the status of the health-related alerts for the objects it is configured to monitor in vRealize Operations Manager. Health alerts usually require immediate attention. You can create one or more health widgets for different objects that you add to your custom dashboards.

How the Health Widget and Configuration Options Work

You can add the Health widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A trend line displays the health status of the monitored object if the object does not provide its resources to any other object. For example, if the monitored object is a virtual machine or a distributed switch.

- A weather map displays the health of the ancestor and descendant objects of the monitored object for all other object types. For example, if the monitored object is a host that provides CPU and memory to a virtual machine.

If the **Badge Mode** is set to **On**, only the badge appears.

You edit a Health widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Health Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Heath Widget Display Options

The Health widget displays a health badge. The widget also displays a health trend when not in badge mode.

Option	Description
Health Badge	Status of the objects configured for this instance of the widget.
	Click the badge to open the **Alerts** tab for the object that provides data to the widget.
	If the **Badge Mode** option is off, a health weather map or trend chart appears for the object. Whether the map or chart appears depends on the object type. The health weather map displays tool tips for up to1000 objects.
Health Trend	Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears.

Heath Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Badge Mode	Determines whether the widget displays only the badge, or the badge and a weather map or trend chart. Select one of the following options: ■ On. Only the badge appears in the widget. ■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Health Chart Widget

The Health Chart widget displays Health, Risk, Efficiency, or custom metric charts for selected objects. You use the widget to compare the status of similar objects based on the same value or name.

How the Health Chart Widget and Configuration Options Work

You can add the Health Chart widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

If the widget is configured to display Health, Risk, or Efficiency, the chart values are based on the generated alerts for the selected alert type for the selected objects.

If the widget is configured to display custom metrics, chart values are based on the metric value for the configured time period.

You edit the Health Chart widget after you add it to the dashboard. The changes you make to the options create a custom widget with the selected charts.

The charts are based either on Health, Risk, or Efficiency alert status, or you can base them on a selected metric. You can include a single object, multiple objects, or all objects of a selected type.

To view the value of the object at a particular time, point your cursor over the chart. A date range and metric value tool tip appear.

A context drop-down menu for each chart can be accessed at the top-right corner after the last metric value.

For each chart, you can view the minimum, maximum, and last metric values. The values are displayed at the top-right corner of each chart. Each of the values is preceded by an appropriate icon of the same color as the state of the metric value.

If there is not enough space to view the metric values, a blue information icon is displayed. Point your cursor over the icon to view the metric value details.

Where You Find the Health Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Health Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Date Controls	Use the date selector to limit the data that appears in each chart to the time period you are examining.
	Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.
	Dashboard Time is the default option.

Health Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

Option	Description
Close	Deletes the chart.
Save a snapshot	Creates a PNG file of the current chart. The image is the size that appears on you screen.
	You can retrieve the file in your browser's download folder.
Save a full screen snapshot	Downloads the current graph image as a full-page PNG file, which you can display or save.
	You can retrieve the file in your browser's download folder.
Download comma-separated data	Creates a CSV file that includes the data in the current chart.
	You can retrieve the file in your browser's download folder.
Units	Select the units in which the widget displays data. This option is visible when you select a custom source of data in the widget configuration.

Health Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Order By	Determines how the object charts appear in the widget.
	You can order them based on value or name, and in ascending or descending order.
Chart Height	Controls the height of all charts. Choose from three possible choices - Small, Medium, Large. Default is Medium.
Pagination number	Number of charts that appears on a page.
	If you prefer scrolling through the charts, select a higher number. If you prefer to page through the results, select a lower number.
Auto Select First Row	Determines whether to start with the first row of data.

| Option	Description
Metric	Determines the source of the data. ■ Health, Risk, or Efficiency. The displayed charts are based on one of these alert badges. ■ Custom. The displayed charts are based on the selected metric and use either alert symptom state colors or the selected custom color. You can select a unit for the custom metric from the drop-down menu or choose to allow the widget to automatically pick a unit. If you apply custom colors, enter the value in each box that is the highest or lowest value that should be that color. You can select a unit for the metric.
Metric Unit	Select a unit for the custom metric.
Show	Select one or more of the following items to display in the widget: ■ Select **Object Name** to display the name of the object in the widget. ■ Select **Metric Name** to display the name of the metric in the widget.
Input Data	
Objects	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	

Option	Description
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. If the objects have an input transformation applied, you select tag values for the transformed objects.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types. If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied. If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects. 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click **Add**. 5 To add another filter criteria set, click **Add another criteria set**.

Mashup Chart Widget

The Mashup Chart widget shows disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs).

How the Mashup Chart Widget and Configuration Options Work

The Mashup Chart widget contains charts that show different aspects of the behavior of a selected resource. By default, the charts show data for the past six hours.
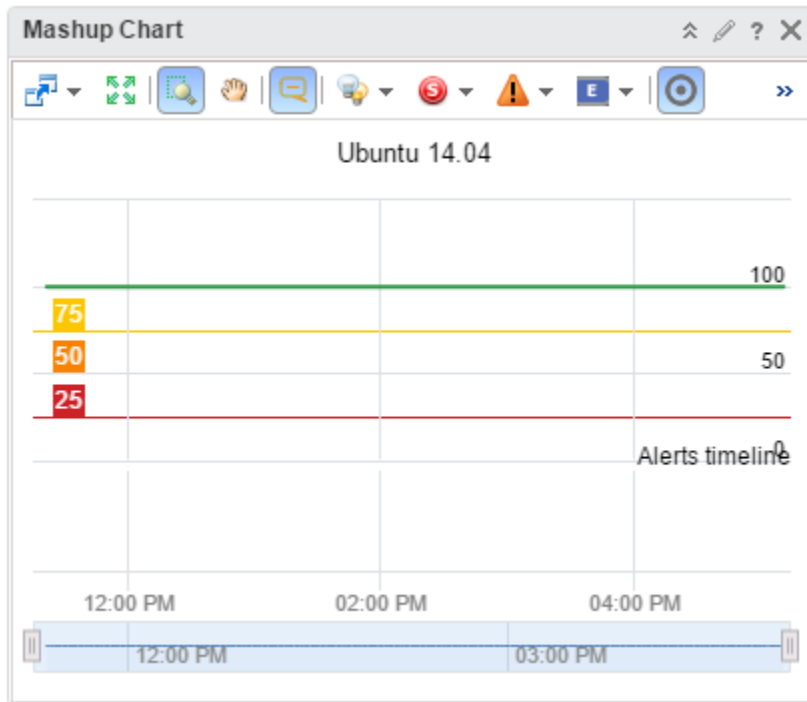
The Mashup Chart widget contains the following charts.

- A Health chart for the object, which can include each alert for the specified time period. Click an alert to see more information, or double-click an alert to open the Alert Summary page.

- Metric graphs for any or all the KPIs for any objects listed as a root cause object. For an application, this chart shows the application and any tiers that contain root causes. You can select the KPI to include by selecting **Chart Controls > KPIs** on the widget toolbar. Any shared area on a graph indicates that the KPI violated its threshold during that time period.

The metric graphs reflect up to five levels of resources, including the selected object and four child levels.

You edit a Mashup Chart widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where You Find the Mashup Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Mashup Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view.

Option	Description
Filters	Filter data based on criticality, status, and alert type.
Event Filters	Filter based on the type of event such as, change, notification, and fault.
Date Controls	Use the date selector to limit the data that appears in each chart to the time period you are examining.
	Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.
	Dashboard Time is the default option.
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate.

Mashup Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Metric Chart Widget

You can use the Metric Chart widget to monitor the workload of your objects over time. The widget displays data based on the metrics that you select.
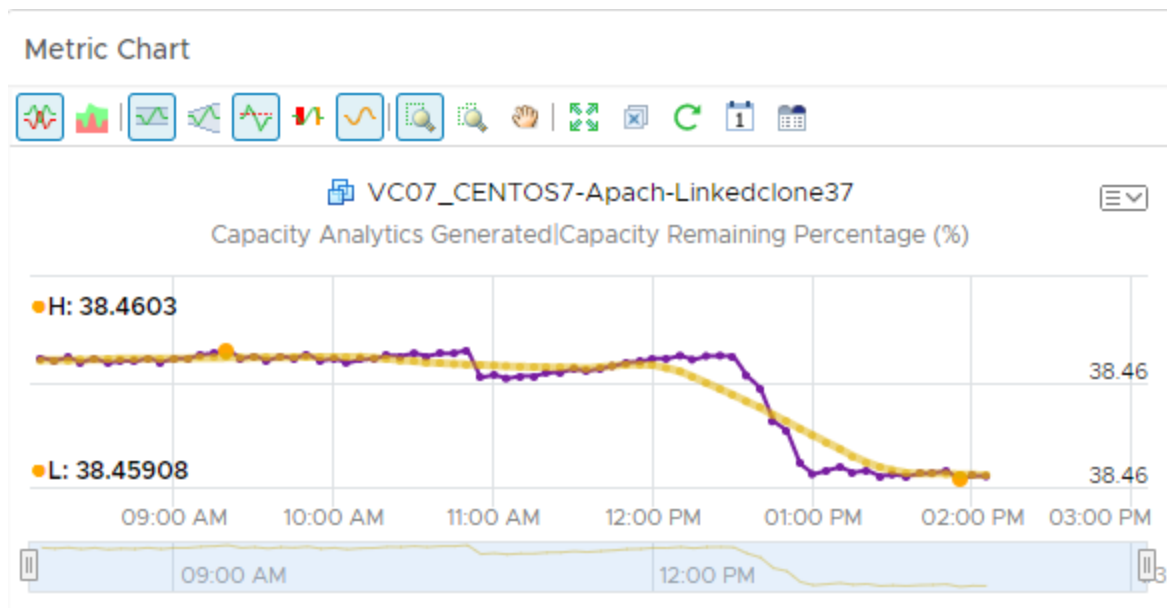
How the Metric Chart Widget and Configuration Options Work

You can add the Metric Chart widget to one or more custom dashboards and configure it to display the workload for your objects. The data that appears in the widget is based on the configured menu items for each widget instance.

You edit the Metric Chart widget after you add it to a dashboard. The changes you make to the menu items create a custom widget with the selected metrics that display the workload on your objects.

To select metrics, you can select an object from the object list, then select the metrics. Or, you can select a tag from the object tag list to limit the object list, then select an object. You can configure multiple charts for the same object or multiple charts for different objects.

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** menu items 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** options are set to **Off**.

- The custom XML file in the **Metric Configuration** drop-down menu is in the `/usr/lib/ vmware-vcops/tools/opscli` directory and has been imported into the global storage using the import command.

Where You Find the Metric Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Metric Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Split Charts	Displays each metric in a separate chart.
Stacked Chart	Consolidates all charts into one chart. This chart is useful for seeing how the total or sum of the metric values vary over time. To view the stacked chart, ensure that the split chart option is turned off.
Dynamic Thresholds	Shows or hides the calculated dynamic threshold values for a 24-hour period.
Show Entire Period Dynamic Thresholds	Shows or hides dynamic thresholds for the entire time period of the graph.
Static Thresholds	Shows or hides the threshold values that have been set for a single metric.
Anomalies	Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below.
Trend Line	Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points.
Show Data Values	Enables the data point tooltips if you switched to a zoom or pan option. **Show Data Point Tips** must be enabled.

Option	Description
Zoom All Charts	Resizes all the charts that are open in the chart pane based on the area captured when you use the range selector. You can switch between this option and **Zoom the View**.
Zoom the View	Resizes the current chart when you use the range selector.
Pan	When you are in zoom mode, allows you to drag the enlarged section of the chart so that you can view higher or lower, earlier or later values for the metric.
Zoom to Fit	Resets the chart to fit in the available space.
Remove All	Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts.
Refresh Charts	Reloads the charts with current data.
Date Controls	Opens the date selector. Use the date selector to limit the data that appears in each chart to the time period you are examining. Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. **Dashboard Time** is the default option.
Generate Dashboard	Saves the current charts as a dashboard.

Metric Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

Option	Description
Close	Deletes the chart.
Save a snapshot	Creates a PNG file of the current chart. The image is the size that appears on you screen. You can retrieve the file in your browser's download folder.
Download comma-separated data	Creates a CSV file that includes the data in the current chart. You can retrieve the file in your browser's download folder.
Save a full screen snapshot	Downloads the current graph image as a full-page PNG file, which you can display or save. You can retrieve the file in your browser's download folder.
Units	You can display the data with dots or as a percentage.
Thresholds	You can choose to show/hide **Critical**, **Immediate**, and **Warning** thresholds in the current chart.

Option	Description
Scales	You can choose a scale for a stacked chart. ■ Select **Linear** to view a chart in which the Y axis scale increases in a linear manner. For example, the Y axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on. ■ Select **Logarithmic** to view a chart in which the Y axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values. **Note** If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph. ■ Select **Combined** to view overlapping graphs for the metrics. The chart uses individual scales for each graph instead of using a relative scale, and displays a combined view of the graphs. ■ Select **Combined by Unit** to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs.
Move Down	Moves the chart down one position.
Move Up	Moves the chart up one position.

You can take the following actions on the Metric Chart graph.

Option	Description
Y Axis	Shows or hides the Y-axis scale.
Chart	Shows or hides the line that connects the data points on the chart.
Data Point Tips	Shows or hides the data point tooltips when you hover the mouse over a data point in the chart.
Zoom by X	Enlarges the selected area on the X axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simultaneously.
Zoom by Y	Enlarges the selected area on the Y axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simultaneously.
Zoom by Dynamic Thresholds	Resizes the Y axis of the chart so that the highest and the lowest values on the axis are the highest and the lowest values of the dynamic threshold calculated for this metric.
Vertical resize	Resizes the height of a graph in the chart.
Remove icon next to each metric name in a stacked chart	Removes the graph for the metric from the chart.

Metric Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	

| Option	Description
Metrics	Select metrics on which you want to base the widget data. You can select an object and pick its metrics.
	1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.
	The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.
	While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.
	Click the **Select All** icon to select all the metrics in the list.
	Click the **Clear Selection** icon to clear your selection of metrics in the list.
	Optionally, you can customize a metric and apply the customization to other metrics in the list.
	1 Double-click a metric box in the list to customize the metric and click **Update**.
	You can use the **Box Label** text box to customize the label of a metric box.
	You can use the **Unit** text box to define a measurement unit of each metric.
	You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.
	For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.
	2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.
Objects	Select objects on which you want to base the widget data.
	1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.

Option	Description
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Empty drop-down menu	Specifies a list with attributes to display.
	To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration. The newly created XML file appears in this drop-down menu.
	Select metrics on which you want to base the widget data. You can select an object and pick its metrics.
	1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.
	The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.
	While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.
	Click the **Select All** icon to select all the metrics in the list.
	Click the **Clear Selection** icon to clear your selection of metrics in the list.
	Optionally, you can customize a metric and apply the customization to other metrics in the list.
	1 Double-click a metric box in the list to customize the metric and click **Update**.
	You can use the **Box Label** text box to customize the label of a metric box.
	You can use the **Unit** text box to define a measurement unit of each metric.
	You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.
	For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.
	2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.

Option	Description
Output Filter	
	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.

1   In the first drop-down menu, select an object type.

2   In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.

3   In the drop-down menus and text boxes that appear, select or enter values to filter the objects.

4   To add more filter criteria, click **Add**.

5   To add another filter criteria set, click **Add another criteria set**.

Metric Picker Widget

The Metric Picker widget displays a list of available metrics for a selected object.
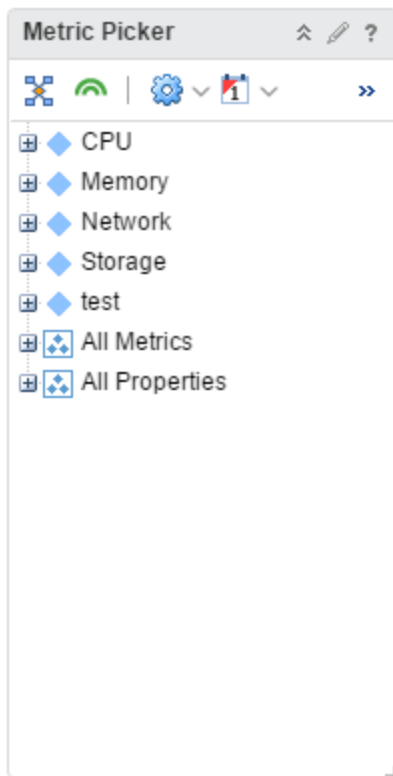
How the Metric Picker Widget and Configuration Options Work

With the Metric Picker widget, you can check the list of the object's metrics. To select an object to pick its metrics, you use another widget as a source of data, for example, Topology Graph widget. To set a source widget that is on the same dashboard, you use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a dashboard that contains the source widget. You can also search for objects using tags.

You edit a Metric Picker widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

Where You Find the Metric Picker Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Metric Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Show common metrics	Filter based on common metrics.
Show collecting metrics	Filter based on collecting metrics.
Metrics or Properties	Filter based on metrics or property metrics.
Time Range	Filter based on selected time range.
Search	Search for dashboards, views, and network IP addresses using tags.

Metric Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Action
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Object List Widget

The Object List widget displays a list of the objects available in the environment.

How the Object List Widget and Configuration Options Work

The Object List widget displays a data grid with objects in the inventory. The default configuration of the data grid appears in Object List Widget Options section. You can customize it by adding or removing default columns. You can use the **Additional Column** option to add metrics when you configure the widget.

You edit an Object List widget after you add it to a dashboard. Configuration of the widget enables you to observe parent and child objects. You can configure the widget to display the child objects of an object selected from another widget, for example, another Object List or Object Relationship widget, in the same dashboard.

Click the legend at the bottom of the widget to filter the objects based on threshold. Point your cursor over any of the boxes to view tooltips.

Where You Find the Object List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Action	Selects from a set of actions specific for each object type. To see available actions, select an object from the list of objects and click the toolbar icon to select an action. For example, when you select a datastore object in the graph, you can select **Delete Unused Snapshots for Datastore**.
Dashboard Navigation	Navigates you to the object. For example, when you select a datastore from the list of objects and click **Dashboard Navigation**, you can open the datastore in vSphere Web Client.
Reset Grid Sort	Returns the list of resources to its original order.
Reset Interaction	Returns the widget to its initial configured state and undoes any interactions selected in a providing widget. Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards.
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Perform Multi-Select Interaction	If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items. Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to enable the interaction.
Display Filtering Criteria	Displays the object information on which this widget is based.
Page Size	
Filter	Locate data in the widget. You can search for objects or filter the list based on the values of the metrics or properties in the additional columns of the Configuration section.

Object List Widget Data Grid Options

The data grid provides a list of inventory objects on which you can sort and search.

Option	Description
ID	Unique ID for each object in the inventory, randomly generated and produced by vRealize Operations Manager.
Name	Name of the object in the inventory.
Description	Displays the short description of the object given during creation of the object

Option	Description
Adapter Type	Shows the adapter type for each object .
Object Type	Displays the type of the object in the inventory.
Policy	Displays policies that are applied to the object. To see policy details and create policy configurations, in the menu click **Administration**, and then in the left pane click **Policies**.
Creation Time	Displays the date, time, and time zone of the creation of an object that was created in the inventory.
Identifier 1	Can contain the custom name of the object in the inventory or default unique identifier, depending on the type of inventory object. For example, My_VM_1 for a VM in the inventory, or 64-bit hexadecimal value for vRealize Operations Manager Node.
Identifier 2	Can contain the abbreviation of an object type and the unique decimal number or parent instance, depending on the type of the object. For example, vm-457 for a VM and an IP address for vRealize Operations Manager Node .
Identifier 3	Can contain a unique number identifying an adapter type. For example, 64-bit hexadecimal value for vCenter Adapter
Identifier 4	Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses.
Identifier 5	Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses.
Object Flag	Displays a badge icon for each object. You can see the status when you point to the badge.
Collection State	Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the state icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**.
Collection Status	Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the status icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**.

Option	Description
Relevance	Displays the user interest on objects based on the number of clicks. The relevance is determined using a system-wide ranking algorithm that rates the object with most clicks as most relevant object.
Internal ID	Unique number that vRealize Operations Manager uses to identify the object internally. For example, the internal ID appears in log files used for troubleshooting.

Object List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Auto Select First Row	Determines whether to start with the first row of data.
Input Data	
Objects	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
	If the objects have an input transformation applied, you select tag values for the transformed objects.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.
Additional Columns	

Option	Description
Empty drop-down menu	Specifies a list with attributes to display. To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration. The newly created XML file appears in this drop-down menu.
	Add metrics based on object types. The selected metrics are displayed as additional columns in the widget. 1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section. While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree. For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center. 2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**.

Object Relationship Widget

The Object Relationship widget displays the hierarchy tree for the selected object. You can create one or more hierarchy trees in vRealize Operations Manager for the selected objects that you add to your custom dashboards.

How the Object Relationship Widget and Configuration Options Work

You can add the Object Relationship widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an Object Relationship widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

Where You Find the Object Relationship Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object Relationship Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To be able to navigate to another dashboard, configure the relevant option when you create or edit the dashboard.
Badge	Displays the Health, Risk, or Efficiency alerts on the objects in the relationship map. You can select a badge for objects that appear in the widget. The tool tip of a badge shows the object name, object type, and the name of the selected badge with the value of the badge. You can only select one badge at a time.
Zoom to fit	Resets the chart to fit in the available space.
Pan	Click this icon and click and drag the hierarchy to show different parts of the hierarchy.
Show values on point	Shows or hides the data point tooltips when you hover the mouse over a data point in the chart.
Zoom the view	Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.
Zoom in	Zooms in on the hierarchy.
Zoom out	Zooms out on the hierarchy.
Reset to Initial Object	If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size.

Option	Description
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Show Alerts	Select the resource in the hierarchy and click this icon to show alerts for the resource. Alerts appear in a pop-up window. You can double-click an alert to view its Alert Summary page.

Object Relationship Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Auto Zoom to Fixed Node Size	You can configure a fixed zoom level for object icons in the widget display. If your widget display contains many objects and you always need to use manual zooming, this feature is useful because you can use it to set the zoom level only once.
Node Size	You can set the fixed zoom level at which the object icons display. Enter the size of the icon in pixels. The widget shows object icons at the pixel size that you configure.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types. If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied. 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click **Add**. 5 To add another filter criteria set, click **Add another criteria set**.

Object Relationship (Advanced) Widget

The Object Relationship (Advanced) widget displays a graph or tree view that depicts the parent-child relationship of the selected object. It provides advanced configuration options. You can create a graph or tree view in vRealize Operations Manager for the selected objects that you add to your custom dashboards.

How the Object Relationship (Advanced) Widget and Configuration Options Work

You can add the **Object Relationship (Advanced)** widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an **Object Relationship (Advanced)** widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

You can double-click any object in the graph or tree view and see the specific parent-child objects for the focus object. When you double-click the object again, you see the original graph or tree view. If you point your cursor over an object icon, you see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object.

Where You Find the Object Relationship (Advanced) Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object Relationship (Advanced) Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Options	Description
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To navigate to another dashboard, configure the relevant option when you create or edit the dashboard.
Reset to Initial Object	If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.
View Tree/View graph	Displays a tree or graph view of the relationships.
Vertical/Horizontal	Displays a vertical or horizontal view of the graph or tree view.

Options	Description
Hide Text/Show Text	Hides or displays the object names.
Standard View/Fit View	The **Standard View** option fixes the view to a specific zoom level The **Fit View** option adjusts the graph or tree view to fit the screen.
Group Items/Ungroup Items	Groups by objects types. You can view further details by double-clicking on the object. You can also choose to display the graph or tree view without grouping the object types.
Path Exploration	Displays the relative relationship path between two selected objects on the graph or tree view. To highlight the path, click the **Path Exploration** icon and then select the two objects from the graph or tree view.
Layers	■ **Parent/Child**: Displays a graph or tree view of the parent and child relationship for the specific object selected. ■ **Custom**: Indicates the relationship between the objects that are part of the custom relationship. These objects have a connection via the selected custom relationship.
Quick Filter	Enter the name of an object that you want to see in the graph or tree view.

Object Relationship (Advanced) Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Name	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Parents Depth	Select the depth of parent objects to be displayed.
Children Depth	Select the depth of child objects to be displayed.
Inventory trees	Select an existing predefined traversal spec for the initial object relationship graph or tree view.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Property List Widget

You can use the Property List widget to view the properties of objects and their values.

How the Property List Widget and Configuration Options Work

To observe the properties of objects in the Property List widget, you can select object property metrics when you configure the widget itself (Self Provider mode enabled). Alternatively, you can select objects or object property metrics from another widget (Self Provider mode disabled). You can also view a default or custom set of properties by selecting a preconfigured XML file in the Metric Configuration drop-down menu of the widget configuration window.

You edit a Property List widget after you add it to a dashboard. You can configure a widget to receive data from another widget by selecting **Off** for Self Provider mode. When the widget is not in Self Provider mode, it displays a set of predefined properties and their values of an object that you select on the source widget. For example, you can select a host on a Topology widget and observe its properties in the Property List widget. To configure the Property List as a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To configure a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

Where You Find the Property List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Property List Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Object Name	Name of the object, whose properties you observe. You can sort the properties by object name. To open the Object Details page, click an object name.
Property Name	Name of the property. You can sort the properties by property name.
Value	Value of the property. You can sort the properties by value.

Property List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Visual Theme	Select a predefined visual style for each instance of the widget. The options are: Original and Compact.
Show Metric Full Name	You can choose to view the full name of the metrics. The options are: On and Off.
Input Data	
Metrics	Select metrics on which you want to base the widget data. You can select an object and pick its metrics. 1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. The metric tree shows common metrics for several objects when you click the **Show common metrics** icon. While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics. Click the **Select All** icon to select all the metrics in the list. Click the **Clear Selection** icon to clear your selection of metrics in the list. You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.

Option	Description
Objects	Select objects on which you want to base the widget data.
	1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Empty drop-down menu	Specifies a list with attributes to display.
	To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration.
	The newly created XML file appears in this drop-down menu.

Option	Description
	1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.
	While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.
	For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.
	2 Optionally, you can define measurement units for the metrics and properties in the list. Double-click a metric or properties box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.
	3 You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.
Output Filter	
	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Recommended Actions Widget

The Recommended Actions widget 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.

How the Recommended Actions Widget and Configuration Options Work

The Recommended Actions widget appears on the Home dashboard, and displays the health status for the objects in your vCenter Server instance. At a glance, you can see how many objects are in a critical state, and how many objects need immediate attention.

From the Recommended Actions widget, you can focus in on problems further by, for example, clicking an object where the alerts triggered, and by clicking an individual alert.

You can edit the Recommended Actions widget on the Home dashboard, or on another dashboard where you add the widget. With the widget configuration options, you can assign a new name to the widget, set the refresh content, and set the refresh interval.

The Recommended Actions widget includes a selection bar, a summary pane, a toolbar for the data grid, and alert information for your objects in a data grid.

Where You Find the Recommended Actions Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Recommended Actions Widget Selection Bar and Summary Pane

Option	Description
Scope	Allows you to select an instance of vCenter Server, and a data center in that instance.
Object tabs	Displays the object types with the number of objects affected in parentheses. You can display the actions for virtual machines, host systems, clusters, vCenter Server instances, and datastores.
Badge	Select the Health, Risk, or Efficiency badge to display alerts on your objects. Health alerts require immediate attention. Risk alerts require attention in the immediate future. Efficiency alerts require your input to reclaim wasted space or to improve the performance of your objects. For each badge, you can view critical, immediate, and warning alerts. ■ Health Status. With the Health badge selected, displays the number of affected objects and a summary of their health based on the alerts that triggered on the object. Lists the objects that have the worst health, and the number of alerts that triggered on each object. ■ Risk Status. With the Risk badge selected, displays the number of affected objects and a summary of their risk based on the alerts that triggered on the object. Lists the objects that have the highest, and the number of alerts that triggered on each object. ■ Efficiency Status. With the Efficiency badge selected, displays the number of affected objects. Lists the objects that have the lowest efficiency based on the alerts that triggered on the object, and the number of alerts that triggered on each object.
Search filter	Narrows the scope of the objects that appear. Enter a character or a number to search and display an object. When a filter is active, the name of the filter appears below the Search filter text box.

Recommended Actions Widget Toolbar Options

The toolbar allows you to address an alert, and to filter the alert list.

Option	Description
Cancel Alert	Cancels the selected alert. You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again.
Suspend	Suspends an alert for a specified number of minutes. You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object. The user who suspends the alert becomes the assigned owner.
All Filters	Narrows the search to one of the available filter types. For example, you can display all alerts that are related to the Compliance Alert Subtype.

Recommended Actions Widget Data Grid Options

The data grid displays the alerts that triggered on your objects. To resolve the problems indicated by the alerts, you can link to the alerts and the objects on which the alerts triggered.

For more information, see Triggered Alerts.

Option	Description
Criticality	Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon. The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.
Actionable	When an alert has an associated action, you can run the action on the object to resolve the alert.
Suggested Fix	Describes the recommendation to resolve the problem. For example, for Compliance alerts, the recommendation instructs you to use the *vSphere Hardening Guide* to resolve the problem. You can find the *vSphere Hardening Guides* at http://www.vmware.com/security/hardening-guides.html. You can view other available recommendations and their associated actions, if any, to resolve the problem when you click the drop-down menu.
Name	Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name. Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object.
Alert	Name of the alert definition that generated the alert. Click the alert name to view the alert details tabs where you can begin troubleshooting the alert.

Option	Description
Alert Type	Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/Hypervisor, Hardware, Storage, and Network.
Alert Subtype	Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration.
Time	Date and time that the alert triggered.
Alert ID	Unique identification for the alert. This column is hidden by default.

Recommended Actions Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	▪ On. You define the objects for which data appears in the widget.
	▪ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Risk Widget

The risk widget is the status of the risk-related alerts for the objects it is configured to monitor. Risk alerts in vRealize Operations Manager usually indicate that you should investigate problems in the near future. You can create one or more risk widgets for objects that you add to your custom dashboards.

How the Risk Widget and Configuration Options Work

You can add the risk widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the Badge Mode configuration option is set to Off, the badge and a chart appear. The type of chart depends on the object type that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning risk alerts generated over time, if the monitored object is a group.

- A trend line displays the risk status of the monitored object for all other object types.

If the Badge Mode is set to On, only the badge appears.

You edit a risk widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.
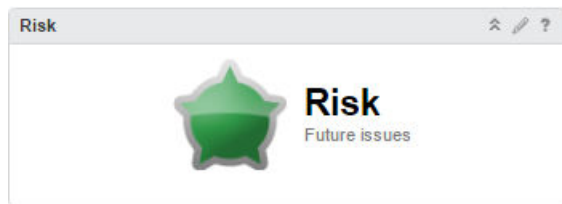
Where You Find the Risk Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Risk Widget Display Options

The Risk Widget displays a risk badge. The widget also displays a risk trend chart when not in badge mode.

Option	Description
Risk Badge	Status of the objects configured for this instance of the widget. Click the badge to open the **Alerts** tab for the object that provides data to the widget.
Risk Trend	Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears.

Risk Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Badge Mode	Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.
	Select one of the following options:
	■ On. Only the badge appears in the widget.
	■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Rolling View Chart Widget

The Rolling View Chart widget 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.
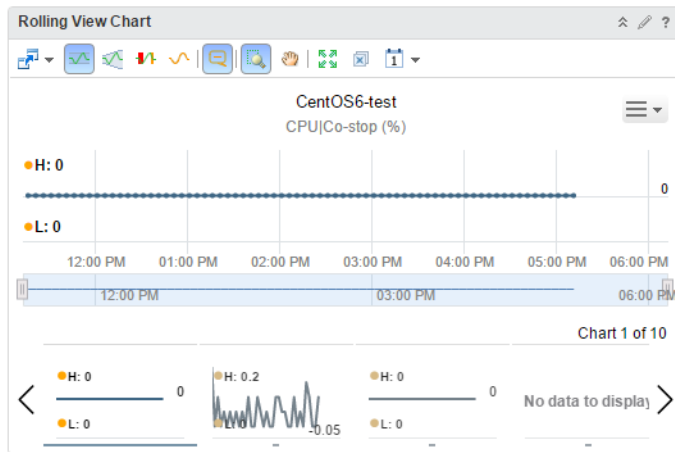
How the Rolling View Chart Widget and Configuration Options Work

The Rolling View Chart widget shows a full chart for one selected metric at a time. Miniature graphs for the other selected metrics appear at the bottom of the widget. You can click a miniature graph to see the full graph for that metric, or set the widget to rotate through all selected metrics at an interval that you define. The key in the graph indicates the maximum and minimum points on the line chart.

You edit a Rolling View Chart widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

Where You Find the Rolling View Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Rolling View Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Trend Line	Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points.
Dynamic Thresholds	Shows or hides the calculated dynamic threshold values for a 24-hour period.
Show Entire Period Dynamic Thresholds	Shows or hides dynamic thresholds for the entire time period of the graph.
Anomalies	Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below.
Zoom to Fit	Changes all graphs to show the entire time period and value range.
Zoom the view	Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section.
Pan	Click this icon and click and drag the hierarchy to show different parts of the hierarchy.
Show Data Values	After you click the **Show data point tips** icon to retrieve the data, click this icon and point to a graphed data point to show its time and exact value. In non-split mode, you can hover over a metric in the legend to show the full metric name, the names of the adapter instances (if any) that provide data for the resource to which the metric belongs, the current value, and the normal range. If the metric is currently alarming, the text color in the legend changes to yellow or red, depending on your color scheme. Click a metric in the legend to highlight the metric in the display. Clicking the metric again toggles its highlighted state.
Date Controls	Use the date selector to limit the data that appears in each chart to the time period you are examining. Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. **Dashboard Time** is the default option.

Rolling View Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Auto Transition Interval	Time interval for a switch between charts in the widget.
Input Data	

Option	Description
Metrics	Select metrics on which you want to base the widget data. You can select an object and pick its metrics. 1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. The metric tree shows common metrics for several objects when you click the **Show common metrics** icon. While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics. Click the **Select All** icon to select all the metrics in the list. Click the **Clear Selection** icon to clear your selection of metrics in the list. You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.

Option	Description
Objects	Select objects on which you want to base the widget data.
	1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Empty drop-down menu	Specifies a list with attributes to display.
	To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration.
	The newly created XML file appears in this drop-down menu.

Option	Description
	Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.
	1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.
	While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.
	For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.
	2 Optionally, you can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.
Output Filter	
	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Scoreboard Widget

The Scoreboard widget shows the current value for each metric of objects that you select.

How the Scoreboard Widget and Configuration Options Work

Each metric appears in a separate box. The value of the metric determines the color of the box. You define the ranges for each color when you edit the widget. You can customize the widget to use a sparkline chart to show the trend of changes of each metric. If you point to a box, the widget shows the source object and metric data. Icons in the box indicate the level of criticality.

You edit a Scoreboard widget after you add it to a dashboard. The widget can display metrics of the objects selected during editing of the widget or selected on another widget. When the Scoreboard widget is not in Self Provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. It shows 10 predefined metrics if you do not select an XML file or if the type of the selected object is not defined in the XML file.

For example, you can configure the Scoreboard widget to use the sample Scoreboard metric configuration and to receive objects from the Topology Graph widget. When you select a host on a Topology Graph widget, the Scoreboard widget shows the workload, memory, and CPU usage of the host.

To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

Where You Find the Scoreboard Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Scoreboard Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
	When the Scoreboard widget is not in self-provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration.
Round Decimals	Select the number of decimal places to round the scores that the widget displays.
Box Columns	Select the number of columns that appear in the widget.
Layout Mode	Select a Fixed Size or Fixed View layout.
Fixed Size Fixed View	Use these options to customize the size of the box for each object.
Old metric values	Select whether you want to show or hide old metric values.
Visual Theme	Select a predefined visual style for each instance of the widget.
Max Scores Count	Use these menus to customize the format of the scores that the widget displays.

Option	Description
Show	Select one or more of the following items to display in the widget: ■ Select **Object Name** to display the name of the object in the widget. ■ Select **Metric Name** to display the name of the metric in the widget. ■ Select **Metric Unit** to display the metric unit in the widget. ■ Select **Sparkline** to display the Sparkline chart for each metric.
Period Length	Select a length of time for the statistic information that the sparkline chart displays.
Show DT	Select an option to show or hide the dynamic threshold for the sparkline chart.

Input Data

Option	Description
Metrics	Select metrics on which you want to base the widget data. You can select an object and pick its metrics. 1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. The metric tree shows common metrics for several objects when you click the **Show common metrics** icon. While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics. Click the **Select All** icon to select all the metrics in the list. Click the **Clear Selection** icon to clear your selection of metrics in the list. Optionally, you can customize a metric and apply the customization to other metrics in the list. 1. Double-click a metric box in the list to customize the metric and click **Update**. You can use the **Box Label** text box to customize the label of a metric box. You can use the **Unit** text box to define a measurement unit of each metric. You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**. For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when

Option	Description
	you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**. You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links open in a new tab. Examples of external links are URLs whose hostname does not match with the current vRealize Operations Manager instance hostname. Internal links are URLs whose hostname matches the current vRealize Operations Manager instance hostname or starts with *index.action*. 2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.
Objects	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	

Option	Description
Empty drop-down menu	Specifies a list with attributes to display.
	To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration.
	The newly created XML file appears in this drop-down menu.

Option	Description
	Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.
	1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.
	While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.
	For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.
	2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.
	Click the **Select All** icon to select all the metrics in the list.
	Click the **Clear Selection** icon to clear your selection of metrics in the list.
	Optionally, you can customize a metric and apply the customization to other metrics in the list.
	1 Double-click a metric box in the list to customize the metric and click **Update**.
	You can use the **Box Label** text box to customize the label of a metric box.
	You can use the **Unit** text box to define a measurement unit of each metric.
	You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.
	For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.

Option	Description
	You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links will open in a new tab. Examples of external links are URLs whose hostname does not match with the current vRealize Operations Manager instance hostname. Internal links are URLs whose hostname matches the current vRealize Operations Manager instance hostname or starts with *index.action*.
	2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.
Output Filter	
	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Scoreboard Health Widget

The Scoreboard Health widget displays color-coded health, risk, efficiency, and custom metrics scores for objects that you select.

How the Scoreboard Health Widget and Configuration Options Work

The icons for each object are color coded to give a quick indication of the state of the object. You can configure the widget to display the scores of common or specific metrics of the object. You can use the symptom state color code or you can define your criteria to color the images. If you configure the widget to show the metric for objects that do not have this metric, those objects have blue icons.

You can double-click an object icon to show the Object Detail page for the object. When you point to the icon, a tool tip shows the name of the object and the name of the metric.
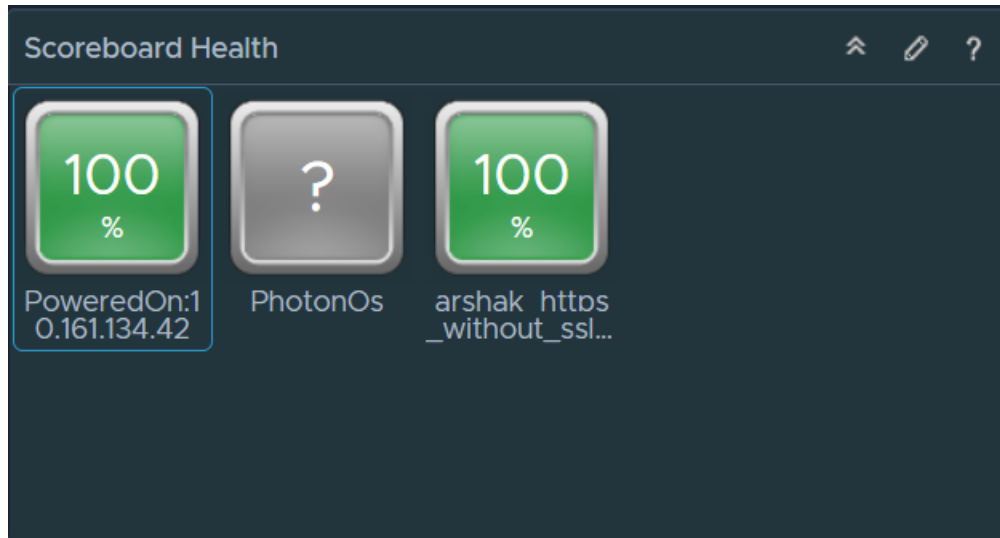
You edit a Scoreboard Health widget after you add it to a dashboard. To configure the widget, click the pencil at the upper-right corner of the widget window. The widget can display metrics of the objects that you select when you edit the widget, or that you select on another widget. For example, you can configure the widget to show the CPU workload of an object that you select on the Topology Graph widget. To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

Where You Find the Scoreboard Health Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Scoreboard Health Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Image Type	Select an image type for the metrics.
Metric	Select the default or custom metric.
Pick Metric	Active only when you select **Custom** from the **Metric** menu. Use to select a custom metric for the objects that the widget displays. Click **Pick Metric** and select an object type from the Object Type pane. Use the Metric Picker pane to select a metric from the metric tree and click **Select Object** to check the objects from the type that you select on the Object Types pane.
Use Symptom state to color chart	Select to use the default criteria to color the image.
Custom ranges	Use to define custom criteria to color the image. You can define a range for each color.

Option	Description
Input Data	
	Select objects on which you want to base the widget data.
	1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.
	While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.
	Click the **Select All** icon to select all the objects in the list.
	Click the **Clear Selection** icon to clear your selection of objects in the list.

Sparkline Chart Widget

The Sparkline Chart widget displays graphs that contain metrics for an object in vRealize Operations Manager. You can use vRealize Operations Manager to create one or more graphs that contain metrics for objects that you add to your custom dashboards.
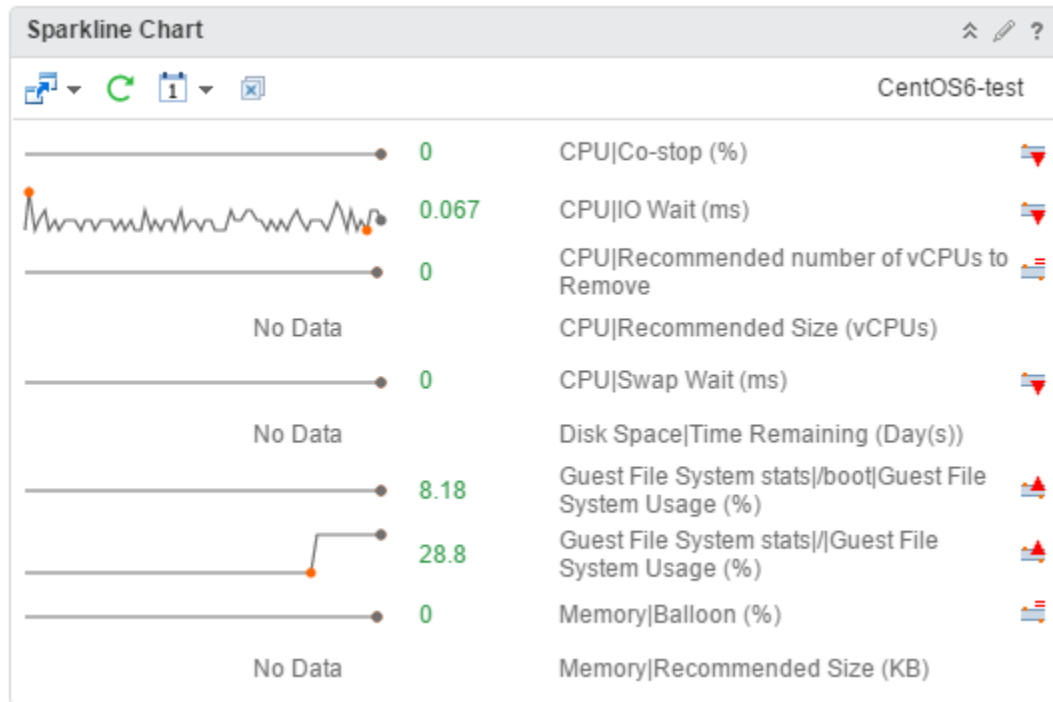
How the Sparkline Chart Widget and Configurations Options Work

If the metrics in the Sparkline Chart are for an object that another widget provides, the object name appears at the top right of the widget. If you select a metric when you edit the widget configuration, the widget uses the metric and its corresponding object as the source for dashboard interactions. The line in the graphs represents the average value of the selected metric for the specified time period. The boxed area in the graph represents the dynamic threshold of the metric.

Point to a graph in the Sparkline Chart widget to view the value of a metric in the form of a tool tip. You can also view the maximum and minimum values on a graph. The values are displayed as orange dots.

You can add the Sparkline Chart widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

Where You Find the Sparkline Chart Widget

The widget might be included on any of your custom dashboards. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Sparkline Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Dashboard Navigation	You can navigate to another dashboard when the object you select is also available in the dashboard to which you want to navigate.
Refresh	Refreshes the widget data.
Time Range	Select the range for the time period to show on the graphs. You can select a period from the default time range list or select start and end dates and times. Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. **Dashboard Time** is the default option.
Remove All	Removes all graphs.

Sparkline Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Show Object Name	You can view the name of the object before the metric name in the Sparkline Chart widget. ■ **On**. Displays the name of the object before the metric name in the widget. ■ **Off**. Does not display the name of the object in the widget.
Column Sequence	Select the order in which to display the information. ■ **Graph First**. The metric graph appears in the first column in the widget display. ■ **Label First**. The metric label appears in the first column in the widget display.
Show DT	Select an option to show or hide the dynamic threshold for the sparkline chart.
Input Data	

Option	Description
Metrics	Select metrics on which you want to base the widget data. You can select an object and pick its metrics.
	1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.
	The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.
	While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.
	Click the **Select All** icon to select all the metrics in the list.
	Click the **Clear Selection** icon to clear your selection of metrics in the list.
	Optionally, you can customize a metric and apply the customization to other metrics in the list.
	1 Double-click a metric box in the list to customize the metric and click **Update**.
	You can use the **Box Label** text box to customize the label of a metric box.
	You can use the **Unit** text box to define a measurement unit of each metric.
	You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.
	For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when

Option	Description
	you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**. 2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.
Objects	Select objects on which you want to base the widget data. 1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
Empty drop-down menu	Specifies a list with attributes to display. To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration. The newly created XML file appears in this drop-down menu.

Option	Description
	Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.
	Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.
	The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.
	While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.
	Optionally, you can customize a metric and apply the customization to other metrics in the list.
	1 Double-click a metric box in the list to customize the metric and click **Update**. You can use the **Box Label** text box to customize the label of a metric box. You can use the **Unit** text box to define a measurement unit of each metric. You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**. For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**. 2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list.

Option	Description
Output Filter	
	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.

Tag Picker Widget

The Tag Picker widget lists all available object tags.

How the Tag Picker Widget and Configuration Options Work

With the Tag Picker widget, you can check the list of the object tags. You can use the widget to filter the information that another widget shows. You can select one or more tags from the object tree or search for tags, and the destination widget displays information about the objects with this tag. For example, you can select **Object Types > Virtual Machine** on the Tag Picker widget to observe statistic information about the VMs on the Environment Status widget.
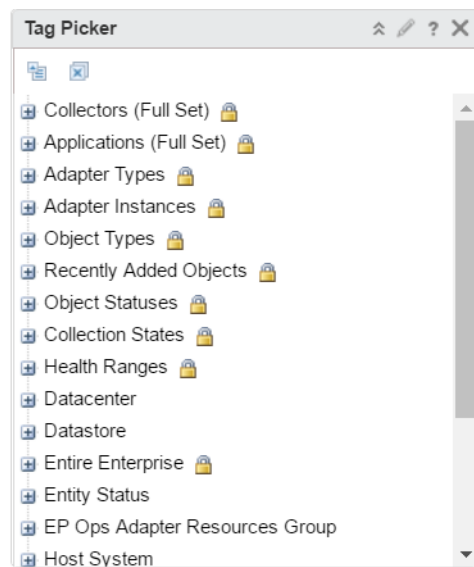
You edit a Tag Picker widget after you add it to a dashboard. To configure the widget, click the pencil in the upper right of the widget window. You can configure the Tag Picker widget to send information to another widget on the same dashboard or on another dashboard. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard. You can configure two Tag Picker widgets to interact when they are on different dashboards.

Where You Find the Tag Picker Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Tag Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Collapse All	Close all expanded tags and tag values.
Deselect All	Remove all filtering and view all objects in the widget.
Tag Picker	Select an object from your environment.
Dashboard Navigation	**Note** Appears on the source widget and when the destination widget is on another dashboard.
	Use to explore the information on another dashboard.

Tag Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Text Display Widget

You can use the Text Display widget to show text in the user interface. The text appears in the Text Display widget on the dashboard.

The Text Display widget can read text from a Web page or text file. You specify the URL of the Web page or the name of the text file when you configure the Text widget. To use the Text Display widget to read text files you must set a property in the *web.properties* file to specify the root folder that contains the file.

You can enter content in the Text Display widget in plain text or rich text format based on the view mode that you configure. Configure the Text Display widget in HTML view mode to display content in rich text format. Configure the Text Display widget in Text mode to display content in plain text format.

The Text Display widget can display websites that use the HTTPS protocol. The behavior of the Text Display widget with websites that use HTTP, depends on the individual settings of the websites.

Note If the webpage that you are linking to has `X-Frame-Options` set to `sameorigin`, which denies rendering a page in an iframe, the Text Display widget cannot display the contents of the webpage.

How the Text Display Widget Configuration Options Work

You can configure the widget in the Text view mode or HTML view mode. In the HTML view mode, you can click **Edit** in the widget and use the rich text editor to add content.

If you configure the widget to use Text view mode, you can specify the path to the directory that contains the files to read or you can provide a URL. The content in the URL will be shown as text. If you do not specify a URL or text file, you can add content in the widget. Double-click the widget and enter content in plain text.

You can also use command-line interface (CLI) commands to add file content to the Text Display widget.

- To view a list of parameters, run the `file -h|import|export|delete|list txtwidget` command.

- To import text or HTML content, run the `import txtwidget input-file [--title title] [--force]` command.

- To export the content to the file, run the `export txtwidget all|title[{,title}] [output-dir]` command.

- To delete imported content, run the `delete txtwidget all|title[{,title}]` command.

- To view the titles of the content, run the `list txtwidget` command.

Where You Find the Text Display Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Text Display Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
View mode	Display text in text or rich text format. You can configure the widget in HTML view mode only when the **URL** and **File** fields are blank.
URL	Enter the URL.
File	Navigate to the file that contains the source text file by clicking the **Browse** button. To add, edit, and remove source text files, go to the **TxtWidgetContent** node in the Metric Configurations page. In the menu, click **Administration**, and then in the left pane click **Configuration > Metric Configurations** from the vRealize Operations Manager user interface.
Test	Validates the correctness of the text file or URL that you enter.

Time Remaining Widget

The Time Remaining widget displays how much time remains before the resources of the object are exhausted.

vRealize Operations Manager calculates the percentage by object type based on historical data for the pattern of use for the object type. You can use the time remaining percentage to plan provisioning of physical or virtual resources for the object or rebalance the workload in your virtual infrastructure.
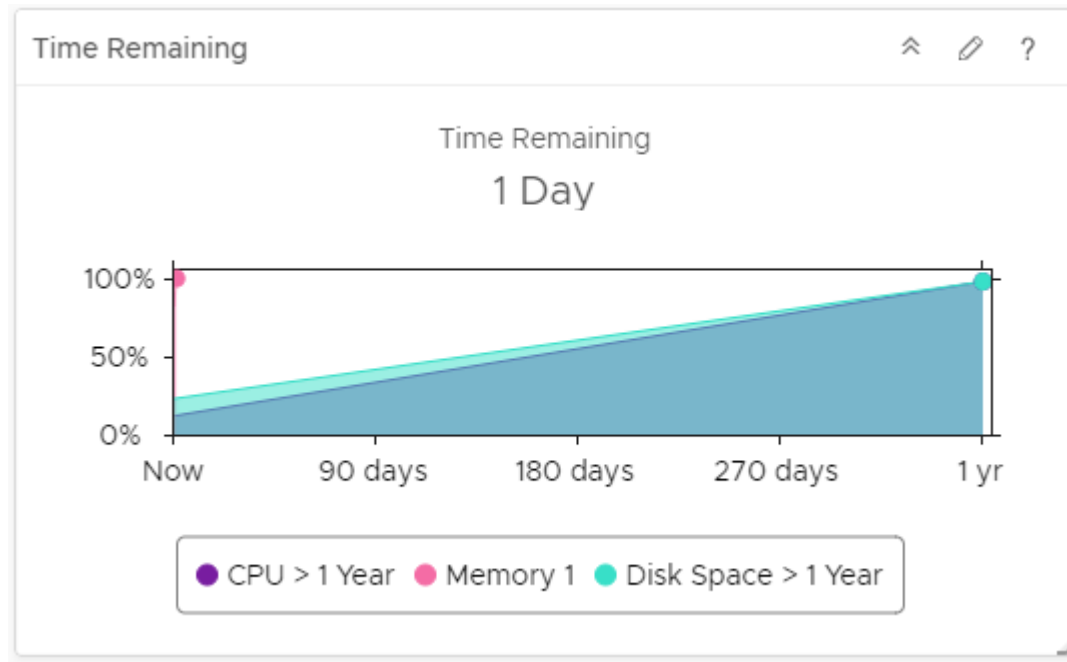
Where You Find the Time Remaining Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Time Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Top Alerts Widget

Top alerts are the alerts with the greatest significance on the objects it is configured to monitor in vRealize Operations Manager. These are the alerts most likely to negatively affect your environment and you should evaluate and address them.

How the Top Alerts Widget and Configuration Options Work

You can add the top alerts widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a top alerts widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.
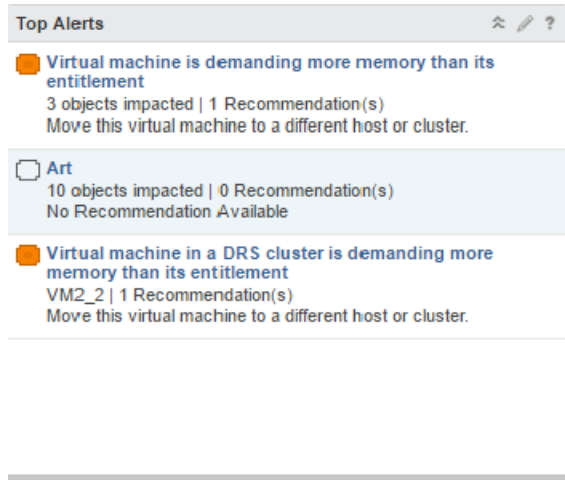
Where You Find the Top Alerts Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Top Alerts Widget Display Options

The Top Alerts widget includes the 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.

Option	Description
Alert name	Name of the generated alert. Click the name to open the alert details.
Alert description	Number of affected objects, and the number of recommendations and the best recommendation to resolve the alert.

Top Alerts Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

| Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Impact Badge	Select the badge for which you want alerts to appear. The affected badge is configured when you configure the alert definition.
Number of Alerts	Select the maximum number of alerts to display in the widget.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.

Top-N Widget

The Top-N widget displays the top n results from analysis of an object or objects that you select.

How the Top-N Widget and Configuration Options Work

You can select an object when you configure the Top-N widget or you can select an object on another widget. The widget shows an analysis of the applications, alerts, and metrics of an object and its child objects depending on how you configure the widget. The widget can show an analysis of the current values or values over a period of time. You can receive detailed information about each object on the widget. When you double-click an object, the Object Detail page appears.

You can configure a widget to receive data from another widget by selecting **Off** for Self Provider. You can configure a widget to display results from analysis of an object that you select on the source widget.

For example, you can select a host on a Topology widget and observe the metric analysis of the virtual machines on the host. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

Where You Find the Top-N Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Top-N Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Icon	Description
Dashboard Navigation	Takes you to a predefined object. For example, when you select a datastore from the data grid and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client.
Select Date Range	Limits the alerts that appear in the list to the selected date range. Select **Dashboard Time** to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.
Object details	Select an object and click this icon to show the Object Detail page for the object.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.

Top-N Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Redraw Rate	Set the redraw rate.
Bars Count	Select the number of top results.
Round Decimals	Select the number of decimals to round the scores displayed in the widget.
Filter old metrics	Select or deselect whether the analysis includes old metric values.

Option	Description
Application Health and Performance	- Top Least Healthy. The top n results from an analysis of the object or objects that are the least healthy. - Top Most Healthy. The top n results from an analysis of the object or objects that are the most healthy. - Top Most Volatile. The sorted list of values based on the standard deviation of values for several alerts over time. Select the criteria for analysis of the objects.
Alert Analysis	Select the criteria for analysis of the alerts.
Metric Analysis	If you select this option, you must select a metric in the **Output Data** section. - Top Highest Utilization. A list of objects with similar object types that have the highest utilization on configuring usage metrics like CPU usage and memory usage. - Top Lowest Utilization. A list of objects with similar object types that have the lowest utilization on configuring usage metrics like CPU usage and memory usage. - Top Abnormal States. The objects are ordered by the duration of all alarms that are triggered on the selected metric for a selected interval. - Top Highest Volatility. The sorted list of values based on the standard deviation of values for several alerts over time. Select the criteria for analysis of the metric that you select from the metric tree.
Input Data	

Option	Description
Objects	Select objects on which you want to base the widget data. 1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears.. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. 2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. Click the **Select All** icon to select all the objects in the list. Click the **Clear Selection** icon to clear your selection of objects in the list.
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget.
Output Data	
	Select an object type in your environment on which you want to base the widget data. 1. Click the **Add Object Type** icon to search for and add an object type. When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. 2. Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type. If the objects have an input transformation applied, the transformed objects are the basis for the widget data.
Metric	Select a common metric or a metric for the selected object type in the list. The metric is the basis for the widget data.

Option	Description
Label	Type in a name that displays as a label for the metric.
	You can add a label if you have selected **Metric Analysis > Top Highest Utilization** or **Metric Analysis > Top Lowest Utilization** as Top-N options in the **Configuration** section.
Unit	You can define measurement units for the metrics. Select a measurement unit in the **Unit** drop-down menu.
	You can add a unit if you have selected **Metric Analysis > Top Highest Utilization** or **Metric Analysis > Top Lowest Utilization** as Top-N options in the **Configuration** section.
Color Method	You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. If you do not want to use color, select None.
	You can add color thresholds if you have selected **Metric Analysis > Top Highest Utilization**, **Metric Analysis > Top Lowest Utilization**, or **Metric Analysis > Percentile** as Top-N options in the **Configuration** section.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
	If the objects have an input transformation applied, you select tag values for the transformed objects.

Option	Description
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.
	If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.
	If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.
	1 In the first drop-down menu, select an object type.
	2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.
	3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.
	4 To add more filter criteria, click **Add**.
	5 To add another filter criteria set, click **Add another criteria set**.
Additional Columns	
	Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.
	1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.
	While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.
	For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.
	2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**.

Topology Graph Widget

The Topology Graph widget gives a graphical presentation of objects and their relationships in the inventory. You can customize each instance of the widget in your dashboard.
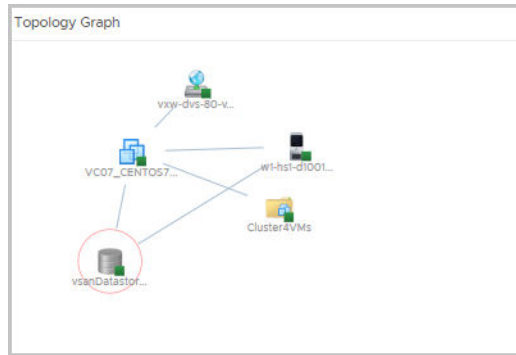
Topology Graph

How the Topology Graph Widget and Configuration Options Work

The Topology Graph widget enables you to explore all nodes and paths connected to an object from your inventory. Connection between the objects might be a logical, physical, or network connection. The widget can display a graph that shows all of the nodes in the path between two objects, or that shows the objects related to a node in your inventory. You select the type of graph in the Exploration Mode when you configure the widget. You can select the levels of exploration between nodes in the displayed graph by using **Relationship** check boxes when you edit the widget. The widget displays all object types in the inventory by default, but you can select object types to view by using the Object View list during the configuration process. Double-clicking an object on the graph takes you to a detailed page about the object.

Where You Find the Topology Graph Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Topology Graph Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Action	Use to select from predefined actions for each object type. To see available predefined actions, select an object in the graph and click the toolbar to select an action. For example, when you select a datastore object in the graph, you can click **Delete Unused Snapshots for Datastore** to apply this action to the object.
Dashboard Navigation	Takes you to a predefined object . For example, when you select a datastore from the graph and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client .

Option	Description
Pan	Use to move the entire graph.
Show values on point	Provides a tool tip with parameters when you point to an object in the graph.
Zoom in	Zooms in the graph.
Zoom out	Zooms out the graph.
Hierarchical View	Use to switch to hierarchical view. Hierarchical view is enabled only for Node Exploration mode and with selected inventory tree.
Graph View	Use to switch to graph view.
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Expand Node	Selects which object types related to your object to show on the graph. For example, if you select a virtual machine from the graph and click **Expand Node** toolbar icon and select **Host System**, the host on which the virtual machine is located is added to the graph.
Hide Node(s)	Use to remove a given object from the graph
Reset To Initial Object	Use to return to the initially displayed graph and configured object types.
Explore Node	Use to explore a node from a selected object in the graph. For example, if the graph displays a connection between a VM, a host, and a datastore, and you want to check the connection of the host with the other objects in the inventory, you can select the host and click **Explore Node**.
Status	Use to select objects based on their status or their state.

Topology Graph Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Exploration Mode	Use **Node Exploration mode** to observe a selected object from an object list and the objects related to it. For example, if you select a virtual machine and select node exploration mode, the widget shows the host where the VM is placed and the datastore storing the files of the VM. Use **Path Exploration mode** to observe the relation between two objects. You must select them from the Select First Object list and the Select Second Object list. For example, if you select to explore the path between a VM and a vCenter Server, the graph shows you both objects and all nodes in the path between the VM and server as datastore, datastore cluster, and datacenter . **Important** To select object view is mandatory for the widget to start working in path exploration mode.
Show Paths	Use **All** to observe connections between a node and nodes related to it as well as connections between the nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph shows a VM connected to its datastore and host and the connection between the host and datastore. Use **Discovered Only** to observe directly related nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph will shows the VM connected to its datastore and to its host, but without the connection between the host and datastore .

Option	Description
Configuration File	The default configuration includes parent and child relationship. Drop-down options depend on the installed Solutions. You can add a new type of relationship to the Relationship pane.
Metric Configuration	Specifies a list with attributes to display. To add a resource interaction XML file through the CLI directory, see Add a Resource Interaction XML File. To add a resource interaction XML file through the UI, see Manage Metric Configuration. The newly created XML file appears in the **Metric Configuration** drop-down menu of the widget.
Layout	Select whether you want a graph view or hierarchical view for the topology graph.
Tree type	For a hierarchical layout, select whether you want a tree type view.
Input Data	
Selected object	From the object list, select an object on which you want to base the widget data.
Degree of separation	Available only when node exploration mode is selected. Use to define the levels of exploration in node exploration mode. The lowest degree configuration shows only directly related nodes rather than higher degrees that show the inventory in details.
Select First Object	Available only in path exploration mode. Select the first object from the object list.
Select Second Object	Available only in path exploration mode. Select the second object from the object list.
Object view	Use to select which types of objects to observe in the graph.
Relationship	Select the type of relationship between objects to observe in the graph, respectively the details about your inventory . The common relationships for all objects are parent and child, but the list of relationships can vary depending on added solutions to vRealize Operations Manager.

View Widget

The View widget provides the vRealize Operations Manager view functionality into your dashboard.

How the View Widget and Configuration Options Work

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, supermetrics, properties, alerts, policies, and data from a different perspective.

You can add the View widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. List views can send interactions to other widgets.

Where You Find the View Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

You can export the view as a CSV file for any view type.

View Widget Toolbar Options

The View widget toolbar depends on the displayed view type.

Option	Description
Export as CSV	You can export the view as a CSV file for any view type.
Open in External Application	Ability to link to another application for information about the object. For example, you have a List view with VMs. You can select any VM and select **Open in External Application** to open the VM in vSphere Web Client.
Time Settings	Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image. ■ Relative Date Range. Select a relative date range of data transformation. ■ Specific Date Range. Select a specific date range of data transformation. ■ Absolute Date Range. 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. The units of time available are: Hours, Days, Weeks, Months, and Years. 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. ■ Dashboard Time. Select this option to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.

Option	Description
Items per page	You can set the number of results that appear in the widget. Available for **List** view only.
Roll up interval	The time interval at which the data is rolled up.
Actions	An action on the selected object. Depends on the object type.
Filter	Limits the list to objects for a specific host, data center, and so on. You can drill-down in the hierarchical level. Available for **List**, **Trend**, and **Distribution** types of Views.
Filter by name	Limits the list to objects of a specific name. Available for **List** view only.

View Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Input Data	
Inventory trees	Select an existing predefined traversal spec to pick an object for the widget data.
Object	In self-provider mode, click the **Add Object** icon to select an object from the object list. The object list is displayed based on the inventory tree selection. You can also search for the object in this text box.
Output Data	
	A list of defined views available for the selected object is displayed.
	You can create, edit, delete, clone, export, and import views directly from the View widget configuration options.
	For more information, see Views.
Auto Select First Row	Determines whether to start with the first row of data for list type views.
Show	Select one or more of the following items to display in the widget: ■ To display the list of legends in the widget, select **Legend**. ■ To display the name of the labels in the widget, select **Labels**.

Weather Map Widget

The Weather Map widget provides a graphical display of the changing values of a single metric for multiple resources over time. The widget uses colored icons to represent each value of the metric. Each icon location represents the metric value for particular resources. The color of an icon changes to show changes in the value of the metric.

How the Weather Map Widget and Configuration Options Work

You can add the Weather Map widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.
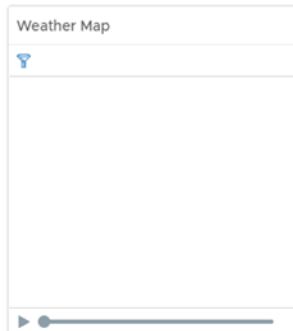
Watching how the map changes can help you understand how the performance of the metric varies over time for different resources. You can start or stop the display using the **Pause** and **Play** options at the bottom of the map. You can move the slider forwards or backwards to a specific frame in the map. If you leave the widget display and return, the slider remains in the same state.

The map does not show the real-time performance of the metrics. You select the time period, how fast the map refreshes, and the interval between readings. For example, you might have the widget play the metric values for the previous day, refreshing every half second, and have each change represent five minute's worth of metric values.

To view the object that an icon represents, click the object.

Where You Find the Weather Map Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Weather Map Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains the icons that you can use to view the graph.

Icon	Description
Pause and **Play**	Start or stop the display. The icon remains in the same state if you leave the widget display and return.
Display Filtering Criteria	View the current settings for the widget, including the current metric.

Weather Map Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Redraw Rate	An interval at which cashed data is refreshed based on newly collected data.
	For example, if you set metric history to `Last 6 hours` and image redraw rate to `15 minutes`, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes.
	For example, if you set metric history to `Last 6 hours` and image redraw rate to `15 minutes`, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes.
Metric History	Select the time period for the weather map, from the previous hour to the last 30 days.
Metric Sample Increment	Select the interval between metric readings. For example, if you set this option to one minute and set the Metric History to one hour, the widget has a total of 60 readings for each metric.
Group by	Select a tag value by which to group the objects.
Sort by	Select **Object name** or **Metric value** to set the way to sort the objects.
Frame Transition Interval	Select how fast the icons change to show each new value. You can select the interval between frames and the number of frames per second (fps).
Start Over Delay	The number of seconds for the display to remain static when it reaches the end of the Metric History period, the most current readings, before it starts over again from the beginning.
Color	Shows the color range for high, intermediate, and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes.
	If you leave the text boxes blank, vRealize Operations Manager maps the highest and lowest values for the **Color By** metric to the end colors.
	If you set a minimum or maximum value, any metric at or beyond that value appears in the end color.

Option	Description
Output Data	
	Select an object type in your environment on which you want to base the widget data.
	1 Click the **Add Object Type** icon to search for and add an object type. When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.
	2 Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type.
Metric	Select a common metric or a metric for the selected object type in the list. The metric will be the basis for the widget data. The object corresponding to the metric is the selected object for the widget.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types. If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied. 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click **Add**. 5 To add another filter criteria set, click **Add another criteria set**.

Workload Widget

The Workload widget displays data indicating how hard a selected resource is working.

The Workload widget displays a graph depicting how hard the object that you selected is working. The Workload widget reports data on CPU usage, Memory usage, Disk I/O, and Network I/O.

Where You Find the Workload Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.
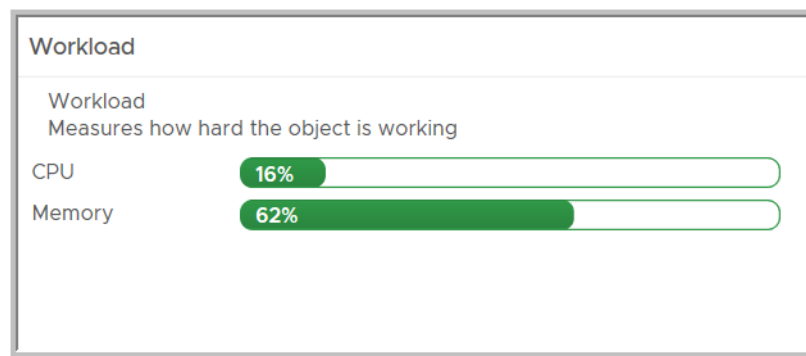
About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.

- The metric named `devices|totalLatency_average` is not supported.

As a result, vRealize Operations Manager does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

Workload Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
	If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.
	■ On. You define the objects for which data appears in the widget.
	■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

Workload Pattern Widget

The Workload Pattern widget displays a historical view of the hourly workload of an object.

Where You Find the Workload Pattern Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Workload Pattern

Workload Pattern
A historical view of hourly workload pattern of an object. This view helps you visualize if an object has been working hard over the last week and identify any hot spots which might cause performance issues.
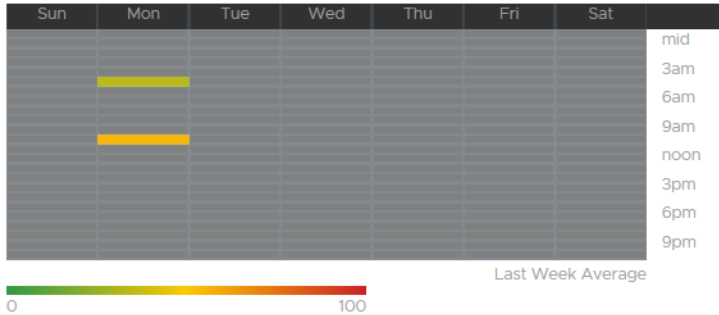
Workload Pattern Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.
Refresh Interval	If you enable the **Refresh Content** option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values.

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 8-4. Menu Options

Menu	Description
Dashboards	Lists the dashboard groups and 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 **Dashboards** option, the dashboards are listed in the left pane of the Dashboards page. The dashboards are listed in the order in which they are selected, with the most recent dashboard selected, appearing at the top. You can reorder the dashboards on the left pane of the **Dashboards** page using drag and drop.
Shared	If the dashboard has been shared, the shared icon is displayed against the dashboard name.

Table 8-4. Menu Options (continued)

Menu	Description
Actions	Available dashboard actions, such as edit, delete, remove dashboard from the menu, set as dashboard landing page, and set as the Home landing page. These actions are applied directly to the dashboard that you are on.
	To remove the dashboard as the Home landing page, from the dashboard that has been set as the Home landing page, select **Actions > Reset from Home landing page**.
	To remove the dashboard as the dashboard landing page, from the dashboard that has been set as the landing page, select **Actions > Reset from Dashboards landing page**.
	You can also create a dashboard and navigate to **Manage Dashboards** page from the **Dashboards** drop-down menu in the left pane.
Dashboard Time	The dashboard time panel is enabled by default on all predefined and user-created dashboards. Using this option, you can select a time for the widgets in the dashboard. The default time is 6 hours. The pre-defined time/day options in the panel are 1 hour, 6 hours, 24 hours, or 7 days. You can also set a customized time option.
	To enable widgets to use the dashboard time, select *Date Controls/Time Range* > **Dashboard Time** from the widget toolbar. Some widgets have **Dashboard Time** as the default option. For example, **Metric Chart**, **View**, **Rolling View**, **Sparkline**, **Health Chart**, and **Mashup Chart** widgets.
	Dashboard time persists if:
	■ You enable a widget in a dashboard to use the dashboard time and then log out and log back in, or
	■ You enable a widget in a dashboard to use the dashboard time, and you export and then import the dashboard into another instance of vRealize Operations Manager.

Types of Dashboards

You can use the predefined dashboards or create your own custom dashboard in vRealize Operations Manager.

See Chapter 9 Predefined Dashboards for more information.

Custom Dashboards

You can create dashboards that meet your environment needs in vRealize Operations Manager.

For information about creating a dashboard, see Create and Configure Dashboards.

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   From the left pane, click the **Dashboards** drop-down menu and then click **Create 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 your dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. Enter a name in the **New Dashboard** field.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

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 `clusters/hosts`, the dashboard is named `hosts` under the group `clusters`.

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 your dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected 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 enter 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.

- To maximize or minimize a widget or a view, use the maximize and minimize options in the top-right corner.

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**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. From the toolbar, click **Show Interactions**.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

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 in the same dashboard or to 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**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a 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.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

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 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 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.

After you have set widget interaction in the provider dashboard, the widget and menu bar are highlighted and two arrows appear in the top-left corner of the widget. After you have set widget interaction, clicking the object in the provider widget takes you to the receiver widget of the navigated dashboard.

Manage Dashboards

You can select dashboards individually or as a group and perform several actions.

To manage your dashboards, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Use the options from the horizontal ellipsis next to the **Add** option.

All the dashboards are listed on this page. You can filter the dashboards based on the name of the dashboard, the dashboard folder, enabled dashboards, shared dashboards, or the dashboard owner. You can click **Add** to create a dashboard. For information about creating a dashboard, see Create and Configure Dashboards.

You can select a dashboard from the list, click the vertical ellipsis against each dashboard, and select the various options such as edit, delete, clone, and disable a dashboard. You can also change ownership of dashboards, save the dashboard as a template, and export the dashboard. By default, the list of dashboards is sorted by name and all the columns can be sorted.

Note A wrench icon appears when the data in an imported dashboard depends on the existence of one or more adapters that are currently not present. The wrench icon disappears if the required data in an imported dashboard appears in vRealize Operations Manager after configuration.

Imported dashboards regardless of used data, remain stuck and include a wrench icon if the dashboard that is stuck (with the wrench icon), already exists.

Datagrid Options

Column Names	Description
Name	Displays the name of the dashboard.
Folder	Lists the folder to which each dashboard belongs.
Description	Displays the description of the dashboard.
Enabled	Enables and disables the dashboard.
URL	Displays whether the dashboard is shared externally. For dashboards that have been shared, click to view the shared links.
Shared	Displays whether the dashboard is shared internally. Click to view and edit the groups to which the dashboard has been shared.
Owner	Displays the owner of the dashboard.
Last Modified	Displays the date the dashboard was last modified.

You can select more than one dashboard and perform a set of options by clicking the horizontal ellipsis next to the **Add** option.

Table 8-5. Dashboards Options

Option	Description	Usage
Export	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.
Enable	Enables a dashboard that was previously disabled.	
Disable	Disables a dashboard.	
Delete	Deletes a dashboard.	
Change Ownership	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.
Import	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.
Auto-rotate 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. For more information, see Auto-Rotate Dashboards.
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. For more information, see Manage Summary Dashboards
Manage Dashboard Folders	Groups dashboards in folders.	You can create dashboard folders to group the dashboards in a way that is meaningful to you. For more information, see Manage Dashboard Folders.
Manage Dashboard Sharing	Makes a dashboard available to other users or user groups.	You can share a dashboard or dashboard template with one or more user groups. For more information, see Share Dashboards with Users.
Clone	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. Accessible as an option only from the vertical ellipsis against the selected dashboard.
Save as Template	Contains all the information in a dashboard definition.	You can use any dashboard to create a template. Accessible as an option only from the vertical ellipsis against the selected dashboard.

The dashboard list depends on your access rights.

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**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and select **Manage Summary Dashboards**.

How You Manage the Summary Dashboards

Table 8-6. Manage Summary Dashboards Toolbar Options

Option	Description
Use Default	Click to use vRealize Operations Manager default **Summary** tab.
Assign a Dashboard	Click to view the Dashboard List dialog box that lists all the available dashboards.
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.

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 All Dashboards dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard that you have associated to the object type when you navigate to the **Summary** tab of the object details page.

Auto-Rotate 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 Auto-Rotation of a Dashboard

To reorder and configure a dashboard switch, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Select **Auto-rotate Dashboards** from the horizontal ellipsis next to the **Add** option.

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 Rotation

1 Double-click a dashboard from the list to configure.

2 From the Rotation drop-down menus, select **On**.

3 Select the time interval in seconds.

4 Select the dashboard to switch 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 Dashboard Folders

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

Where You Manage Dashboard Folders

To manage the dashboard folders, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Folders**.

How You Manage the Dashboard Folders

Table 8-7. Manage Dashboard Folders Options

Option	Description
Dashboards List	A list with all available dashboards.
Folders	A hierarchy tree with all the available group folders.

To create a dashboard folder, click **New Folder** in the **Folders** pane and enter the name of the folder. If you want to create a folder under another folder, select a parent folder under which you want to create the child folder, then click **New Folder**. To add a dashboard, drag one from the dashboards list to the selected folder in the **Folders** pane.

You can delete folders and/or detach dashboards from a folder, by selecting one or more folders and dashboards from the **Folders** pane and by clicking **Actions > Delete**.

You can rename a folder by selecting a single folder from the **Folders** pane and by clicking **Actions > Rename**.

Share Dashboards with Users

You can share a dashboard 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**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Sharing**.

Table 8-8. Dashboard Sharing Options

Option	Description
All Dashboards	Link to view all the available dashboards that you can share. The dashboards are displayed on the right side in the dashboards list.
User Groups	Lists the available user groups that you can share a dashboard with. The list includes the **Everyone** group.
Dashboard List	List of shared dashboards with the selected user group or all the available dashboards that you can share, if no user group is selected.

Manage Dashboard Sharing

To share a dashboard, navigate to the dashboard in the list of 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 **Stop Sharing** above the list.

Dashboards Actions and Options

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.

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.

Dashboard sharing can only be applied to Groups with a vRealize Operations Standard Edition license.

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 8-9. Options in the Share Dashboard Dialog Box

Option	Description
URL	Allows you to copy the tiny URL for the selected dashboard.
	■ 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.
	Note
	■ 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)**.
	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**.
	Authentication is not required to view the shared dashboard.
Email	Allows you to send an email with the URL details of the dashboard, to a specific person.
	■ 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.
	Authentication is not required to view the shared dashboard.

Table 8-9. Options in the Share Dashboard Dialog Box (continued)

Option	Description
Embed	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. ■ Set the expiry period for the link to **1 day**, **1 week**, **1 month**, **3 Months**, or **Never Expire**. **Note** ■ 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. Authentication is not required to view the shared dashboard.
Groups	Allows you to assign and unassign a dashboard to specific user groups. ■ 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. Log in to vRealize Operations Manager to view the shared dashboard.
Export	Allows you to export the dashboard configuration details. Log in to vRealize Operations Manager to export/import a dashboard.

Manage Widgets in Dashboards

You can replicate widgets multiple times in a dashboard by using the copy and paste functionality.

Navigate to the dashboard from which you want to copy widgets. Select **Actions > Edit Dashboard**. Select one or more widgets that you want to copy by clicking the title of the widget and then select **Actions > Copy Widget(s)**. Click **Actions > Paste Widget(s)** to paste one or more widgets in the same dashboard.

To paste one or more widgets into another dashboard, exit the edit screen of the dashboard by selecting **Cancel**. Navigate to the dashboard to which you want to paste one or more widgets and select **Actions > Edit Dashboard** and then **Actions > Paste Widget(s)**.

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. 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.

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. For more information, see View Widget.

- 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.

How You Access the Views Page

In the menu, click **Dashboards**, and then in the left pane click **Views** to access the Views page.

Manage and Preview Views

You can preview a view by clicking a view from the **Views** page. Add an object if necessary, by clicking **Select preview source** from the upper-right corner of the **Views** page. The preview of the view appears just below the **Views** option in the right pane.

You can select a view from the list, click the vertical ellipsis against each view, and select the various options such as edit, delete, clone, and export a view.

You can filter the views based on the name, type, description, subject, and owner. You can click **Add** to create a view. For information about creating a view, see Create and Configure a View.

You can select more than one view and delete, export, and import views by clicking the horizontal ellipsis next to the **Add** option.

Views are also categorized and listed in the **Views** menu based on the type of view and subject. You can access the **Views** menu from a specific view preview page.

Table 8-10. 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 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 **Add** 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**.

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**. Click the **Add** option. In the workspace, on the left, click **Name and Description**.

Table 8-11. 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**. Click the **Add** option. In the workspace, on the left, click **Presentation**. If you create a view, complete the required previous steps.

Table 8-12. 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.
Distribution	Provides aggregated data about resource distribution in the monitored environment.
	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.
Text	Inserts the provided text. The text can be dynamic and contain metrics and properties.
	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.
	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.
Image	Inserts a static image.
	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.

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 8-13. Presentation Configuration Options in the View Workspace

View Type	Configuration Description
List	▪ 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.
Trend	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. 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.
Distribution	Select the visualization of the distribution information in a pie chart or a bar chart. Select the distribution type, and configure the buckets count and size. To understand vRealize Operations Manager distribution type, see View Distribution Type.

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 8-14. 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.

View Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Visualization

You can view the data as a pie chart, a bar chart, or a donut chart. 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, or a section of the donut 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 8-15. 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. You can also select a color for each defined bucket that you specify.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distributes 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**. Click the **Add** option. In the workspace, on the left, click **Subjects**. If you create a 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.

In a List view, you can group the results based on a parent object, by making a selection in the **Group By** drop-down option. If you generate a report based on the list view for which a group has been specified, the report displays group-based information for the selected object. You can also view summary calculations for the group of objects in the report, along with the total summary results for all the objects.

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**. Click the **Add** option. In the workspace, on the left, click **Data**. If you create a 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 8-16. 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.

Table 8-16. Data Configuration Options (continued)

Configuration Option	Description
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 8-16. Data Configuration Options (continued)

Configuration Option	Description
Transformation	Determines what calculation method is applied on the raw data. You can select the type of transformation: ■ 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 Manager6.7, 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 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`. You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps.

Table 8-16. Data Configuration Options (continued)

Configuration Option	Description
	Available for all views, except Trend.
	■ Timestamp: You can choose between `Absolute Timestamp` or `Relative Timestamp`.
	■ If applied to a numeric metric/property defined with a time-unit definition, the actual value is converted to a human readable timestamp. The metric value is rounded-off to an hour.
	Applicable for Absolute Timestamp.
	■ In the remaining cases, a timestamp is displayed when metrics and properties are added or modified. In this case, the behavior is the same as the Timestamp option selected for a non-Timestamp transformation.
	Applicable for Absolute Timestamp and Relative Timestamp.
	Available for List view and Minimum, Maximum, Current, First, and Last transformations.
Ranges for metric coloring	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.
Data Series	You can select whether to include historical data, trend of historical data, and forecast for future time in the trend view calculations.
	Available for Trend view.
Series Roll up	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.
	This option is applicable to the Transformation configuration option.
	Available for all views.
Threshold Lines	You can set a threshold for a single metric:
	■ 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.
	Available for Trend view.

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 8-17. Time Settings Options

Configuration Option	Description
Time Range Mode	In Basic mode, you can select date ranges.
	In Advanced mode, you can select any combination of relative or specific start and end dates.
Relative Date Range	Select a relative date range of data transformation.
	Available in Basic mode.
Specific Date Range	Select a specific date range of data transformation.
	Available in Basic mode.
Absolute Date Range	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.
	The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.
	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.
	Available in Basic mode.
Relative Start Date	Select a relative start date of data transformation.
	Available in Advanced mode.
Relative End Date	Select a relative end date of data transformation.
	Available in Advanced mode.
Specific Start Date	Select a specific start date of data transformation.
	Available in Advanced mode.
Specific End Date	Select a specific end date of data transformation.
	Available in Advanced mode.
Currently selected date range	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`.

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 8-18. Group By Options

Option	Description
Add interval breakdown column (see data for column settings)	Select this option to see the data for the selected resources broken down in time intervals. 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.
Add instance breakdown column (see data for column settings)	Select this option to see the data for all instances of the selected resources. 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. 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`.

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 Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 8-19. Filter Add Options

Option	Description
Add	Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.
	If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.
	For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum.
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 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**. Click the **Add** option. In the workspace, on the left, click **Visibility**. If you create a view, complete the required previous steps.

Table 8-20. 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 data center in the banned list. The view is not visible anymore on data center 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.

Edit a View

When you edit a view, all changes are applied to the report templates that contain it. To edit a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Edit**.

Clone a View

When you clone a view, the changes that you make to the clone do not affect the source view. To clone a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Clone**.

Delete a View

When you delete a view, it is removed from all the report templates that contain it. To delete a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Delete**.

Including Deleted VMs in List View

In vRealize Operations Manager , you can view the deleted objects and the relationship of the objects in the list view. The objects can be VMs, deployments, projects, vApps, and edge gateways. You can also retain the relationship of the objects even after the objects are deleted from the system. The cost of the deleted virtual machines (VMs) is available until the retention period for that VM is over.

Where You Find Global Settings for Deleted VMs

To specify for how long you want to retain the deleted virtual machines in vRealize Operations Manager , navigate to **Administration** > **Management** > **Global Settings** > **Deleted Objects**.

You can also specify the **Deletion Scheduling Interval** which specifies the number of hours between resource deletion scheduling.

In the **Object Deletion Scheduling** page, click **Add**, select the virtual machine object from the drop-down menu, specify the value, and click **Update**. The global setting value for the deleted virtual machine is updated in vRealize Operations Manager .

For vRealize Automation, the price of the deleted VMs or deployments is added to the corresponding project object as a separate metric. If the deleted VM from vRealize Automation is associated with a cost-based pricing policy, then the price for that VM is not added to the corresponding project.

For vCloud Director, the price of deleted VMs, vApps, and Edge Gateways is added to the corresponding organization VDC object again as a separate metric. For vCenter Server, if VM is on unclustered Host, then deleted VM price is assigned to the Host, otherwise to the Cluster.

How to Include Deleted VMs in List View

To view the deleted virtual machines in the list view, navigate to **Dashboards** > **Views** > **Deleted VMs**. Click **Actions** and select **Edit View**. From the filter tab, select **Included Deleted Objects**.

The deleted VMs are visible in **Administration** > **Inventory** > **Object Types** > **Collection States** > **Not Existing**.

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 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 View

To use a view in another vRealize Operations Manager instance, you export a content definition XML file.

4 Import a 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 **Add** option to create a 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 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 View

To use a view in another vRealize Operations Manager instance, 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   Select a view and click **Export** from the vertical ellipsis next to the selected view.

Import a 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 Select a view and click the **Import** option from the horizontal ellipsis.

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.

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, description, subject, date they were last modified, last run, or by whom they were modified.

For more information about the options and actions in the Reports Tab page, see Report Templates Overview.

Table 8-21. Predefined Filter Groups

Filter Group	Description
Name	Filter by the template name. For example, you can list all reports that contain *my template* 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.
Owner	Filter by the owner of the report template.

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.

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.

For more information about the options and actions in the Generated Reports tab page, see Generated Reports Overview.

Table 8-22. Predefined Filter Groups

Filter Group	Description
Report Name	Filter by the report template name. For example, you can list all reports that contain *my template* 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 **Add** 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**.

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**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New Template dialog box, in the workspace, on the left, click **Name and Description**.

Table 8-23. 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**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New 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 8-24. 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.
Quick Filter	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.

Table 8-24. Views and Dashboards Options in the Report Template Workspace (continued)

Option	Description
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.
Colorization	You can enable or disable a colorized PDF output for each list view. This option is available from the right panel when you select **Views** from the **Data type** drop-down menu.

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**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New 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 8-25. 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**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New 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 8-26. 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 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 only SMB version 2.1.

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 **Add**.

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.

Option	Description
Password	The password for the domain user account.
Network share root	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.
	You must enter an IP address. For example, `\\IP_address\ShareRoot`. 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.
	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.

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. See Schedule Reports Overview.

Report Templates Overview

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.

In the menu, click **Dashboards**, and then in the left pane select **Reports**. Select the **Report Templates** tab in the right pane.

The listed templates are user-defined and predefined by vRealize Operations Manager. You can order them by template name, description, subject, date they were modified, last run report, or the user who modified them. For each template, you can see the number of generated reports and schedules.

You can filter the reports based on the name of the report template, the subject, and the owner. You can click **Add** to create a report template. For information about creating a report template, see Create and Modify a Report Template.

You can select a report template from the list, click the vertical ellipsis against each report template, and select options such as run, edit, schedule, delete, clone, and export a report.

Table 8-27. Predefined Filter Groups

Filter Group	Description
Name	Filter by the template name. For example, type `my template` to list all reports that contain the `my template` phrase in their name.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by the other objects.
Owner	Filter by the owner of the report template.

The maximum number of reports per template is 10. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

Report Template Actions

You can select more than one report template and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

Option	Description
Delete	Deletes the report template.
Export	Downloads the report template.
Import	Allows you to import a report template by selecting a report template in XML or zip file format.
Change default cover image	Allows you to change the default cover image of the report template. For more information, see Upload a Default Cover Page Image for Reports.

Generated Reports Overview

A report is a scheduled snapshot of views and dashboards. It presents data in formats that can be downloaded.

In the menu, click **Dashboards**, and then in the left pane select **Reports**. Select the **Generated Reports** tab in the right pane.

The list contains all generated reports. 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. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

To select a generated report from the list, click the vertical ellipsis against each generated report and select options such as run and delete. You can also select more than one generated report and select **Delete** from the **Actions** drop-down menu to delete a generated report.

You can filter the reports list by adding a filter from the upper-right corner of the panel.

Table 8-28. Predefined Filter Groups

Filter Group	Description
Report Name	Filter by the report template name. For example, type `my template` to list all reports that contain the `my template` phrase in their name.
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.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by that second object.
Status	Filter by the status of the report.

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

If you log in to vRealize Operations Manager with vCenter Server credentials and generate a report, the generated report is always blank.

Generate and Regenerate a Report

To generate a report, use a report template.

Prerequisites

Create a report template.

Procedure

1 In the menu, click **Environment**.

2 In the left pane, navigate to the relevant object.

3 Click the **Reports** tab and click **Report Templates**.

   The listed report templates are associated with the current object.

4 Navigate to the relevant report template, click the vertical ellipsis, and select **Run**.

Results

The report is generated and listed on the **Generated Reports** tab.

Note   To regenerate the selected report, from the **Generated Reports** tab, click the vertical ellipsis against the generated report and select **Run**.

What to do next

Download the generated report and verify the output.

Download a Report

To verify that the information appears as expected, you download the generated report.

Prerequisites

Generate a report.

Procedure

1   In the menu, click **Environment**.

2   In the left pane, navigate to the object for which you want to download a report.

3   Click the **Reports** tab and click **Generated Reports**.

The listed reports are generated for the current object.

4   Click the PDF or the CSV icon in the Download column to download the report.

Results

vRealize Operations Manager saves the report file.

What to do next

Schedule a report generation and set the email options, so your team receives the report.

Schedule Reports Overview

The schedule of a report is the time and recurrence of a report generation.

Where Do You Schedule a Report

To schedule a report generation, in the menu, click **Environment**, and then in the left pane navigate to an object and click the **Reports** tab. Select a template to schedule, click the vertical ellipsis, and then click **Schedule**. To edit the schedule of a report, click the **Schedules** link of a report from the **Report Templates** tab, and then from the **Scheduled Reports** dialog box, click **Edit Schedule**.

How Do You Schedule a Report

Table 8-29. Schedule Report Options

Option	Description
Recurrence	Schedule a report to run automatically at regular intervals.
Publishing	Email a generated report to a predefined email group or to a network shared location. For more information about how to set up and configure the email options, see Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.
	Save a generated report to an external location. For more information about how to configure an external location, see Add a Network Share Plug-In for vRealize Operations Manager Reports
	You can add a relative path to upload the report to a predefined sub folder of the Network Share Root folder. For example, to upload the report to the share host `C:/documents/uploadedReports/SubFolder1`, in the **Relative Path** text box, enter `SubFolder1`. To upload the report to the Network Share Root folder, leave the **Relative Path** text box empty.

Note Only users created in vRealize Operations Manager can add and edit report schedules.

Table 8-30. Scheduled Reports Toolbar Options

Options	Description
New Schedule	You can create a schedule for the report.
Edit Schedule	You can edit an existing report schedule.
Delete Schedule	You can delete an existing report schedule.
Transport Report Schedule	You can assign a new owner for the selected report schedule. You can select a target user from the **Transfer Report Schedules** dialog box.

Schedule a Report

To generate a report on a selected date, time, and recurrence, you create a schedule for the report template. You set the email options to send the generated report to your team.

The date range for the generated report is based on the time when generates the report and not on the time when you schedule the report or when vRealize Operations Manager places the report in the queue.

Prerequisites

- Download the generated report to verify the output.

- To enable sending email reports, you must have configured Outbound Alert Settings. See Notifications .

Procedure

1   In the menu, click **Environment**.

2   In the left pane, navigate to the object.

3   Click the **Reports** tab and click **Report Templates**.

4   Select the relevant report template from the list.

5   Click the vertical ellipsis and select **Schedule**.

6   Select the time zone, date, hour, and minutes (in the range of 0, 15, 30, and 45 minutes) to start the report generation.

    vRealize Operations Manager generates the scheduled reports in sequential order. Generating a report can take several hours. This process might delay the start time of a report when the previous report takes an extended period of time.

7   From the **Recurrence** drop-down menu, select one of the following options for report generation:

Option	Description
Daily	You can set the periodicity in days. For example, you can set report generation to every two days.
Weekly	You can set the periodicity in weeks. For example, you can set report generation to every two weeks on Monday.
Monthly	You can set the periodicity in months.

8   Select the **Email report** check box to send an email with the generated report.

    a   In the **Email addresses** text box, enter the email addresses that must receive the report. You can also add email addresses in the CC list and BCC list.

    b   Select an outbound rule.

    An email is sent according to this schedule every time a report is generated.

9   Save a generated report to an external location.

10 You can add a relative path to upload the report to a predefined sub folder of the Network Share Root folder.

    To upload the report to the Network Share Root folder, leave the **Relative Path** text box empty.

11 Click **OK**.

What to do next

You can edit, clone, and delete report templates. Before you do, familiarize yourself with the consequences of these actions.

When you edit a report template and delete it, all reports generated from the original and the edited templates are deleted. When you clone a report template, the changes that you make to the clone do not affect the source template. When you delete a report template, all generated reports are also deleted.

Upload a Default Cover Page Image for Reports

You can upload a common default image for the cover page of reports. You do not have to upload a cover page for each report. The cover pages of predefined reports are modified when you use this option. The cover pages of user-defined reports do not change.

Where Do You Upload a Default Cover Page Image for Reports

To upload a default cover page for reports, in the menu, click **Environment**, and then in the left pane navigate to an object and click the **Reports** tab. From the **Report Templates** tab, click the horizontal ellipsis next to the **Add** option and click the **Change default cover image** option.

How Do You Upload a Default Cover Page Image for Reports

Browse for the image that you want to add to the cover page and click **Save**. You can also use the default product image that is available.

Predefined Dashboards

9

vRealize Operations Manager includes a broad set of simple to use, but customizable dashboards to get you started with monitoring your VMware environment. The predefined dashboards 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.

Each set of dashboards is complemented with a series of out-of-the-box customizable alerts and reports to assist with your operational awareness. Alerts, reports, and dashboards, each have a purpose with minimal overlap. Several activities that are carried out using alerts should be carried out using dashboards. Reports should be kept to a minimum as they are not interactive and do not provide timely information.

The following table details how alerts, dashboards, and reports are complimentary.

	Alerts	Dashboards	Reports
Nature	Reactive	Proactive	Passive. For those with no access to vRealize Operations Manager/vRealize Operations Cloud and vRealize Log Insight.
Suitability	Exception (something went wrong)	Exception Big Picture Details Analysis	Big Picture Exception (but not urgent) No analysis as it is not interactive
Use Case	Troubleshooting (the start)	Monitoring Troubleshooting (the actual)	FYI (optional) Export for further analysis (spreadsheet)
Time & Urgency	Urgent (minutes) & Important	Regular (daily, SOP)	Not urgent (monthly) & optional No point in a daily report. For daily reports, login for interactivity
Access Requirement	Online. Desktop. 1280 * 1024 pixel	Online. Desktop. 1280 * 1024 pixel	Offline or mobile. Small resolution. Email. Laptop or tablet.
Scope/Area	Availability Performance Compliance Configuration (?) Capacity (less relevant, unless it is an emergency)	Availability Performance Capacity Compliance Configuration Inventory	Same as dashboards, but: • without interactivity • time bound (e.g. calendar month) • No performance report, covered in Capacity
Roles	Operations Team	Operations Team Architect Team	IT Management (not hands-on) Auditor (compliance)

Insight vs Alerts

vRealize Operations Manager dashboards support a concept we call insight. Insight complements alerts but does not replace it. Alerts miss the larger picture and only see what is triggered. For one object that reaches the threshold, there might be many just beneath the threshold. The objects below the threshold are called insight.

Alerts might auto-close if the symptoms disappear. Managing alerts is not the same as minimizing alerts. Minimizing alerts is about preventing alerts.
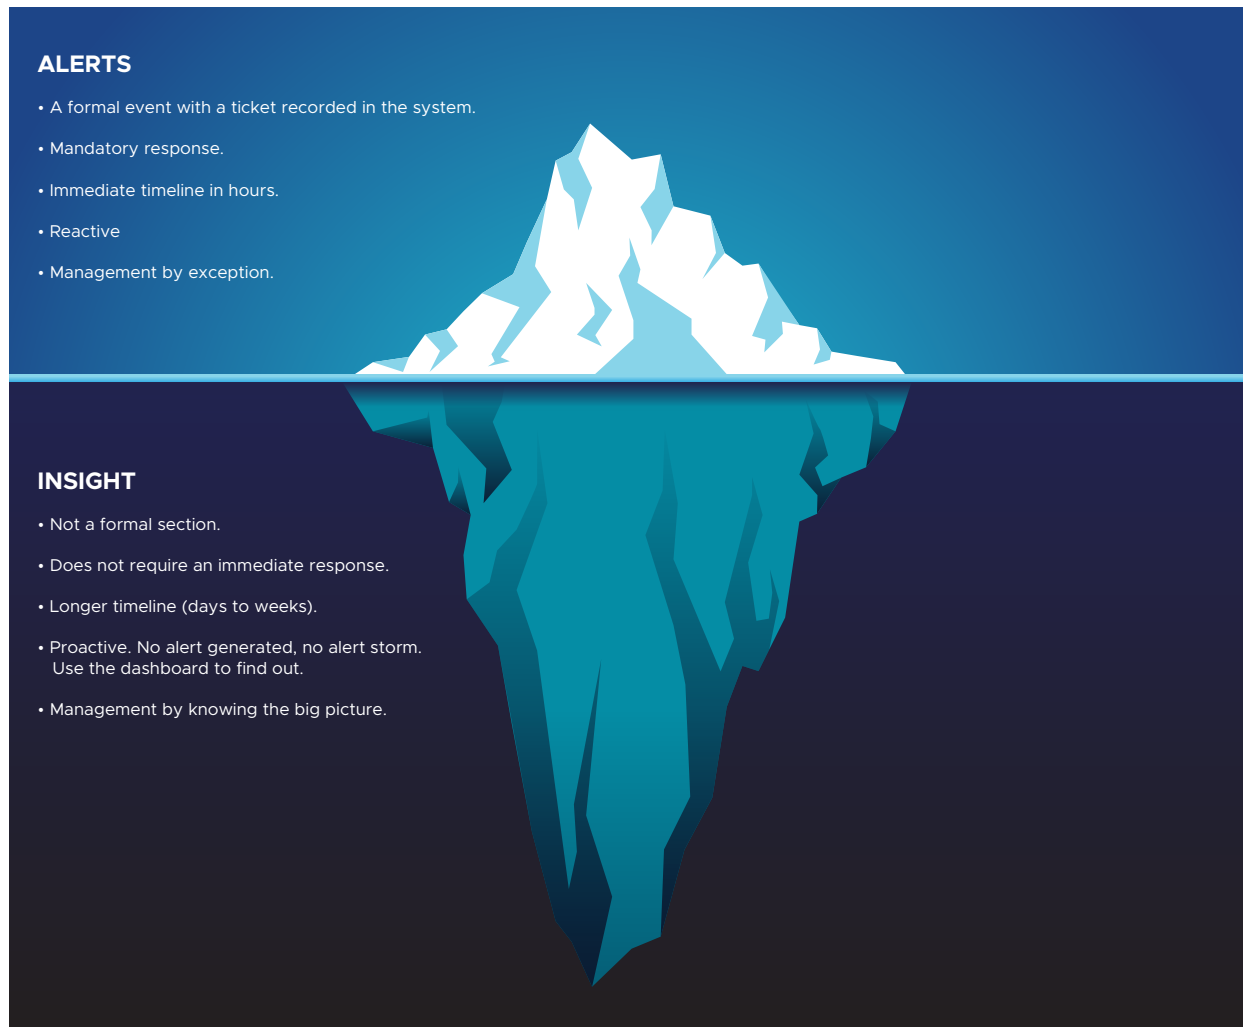
Working with Predefined Dashboards

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.

To access the predefined dashboards, from the left pane, click the **Dashboards** drop-down menu.

To access the deprecated dashboards, from the left pane, click the **Dashboards** drop-down menu and then select **Dashboard Library > Deprecated**.

You can customize dashboards and widgets if you have vRealize Operations Advanced edition or higher. Any customization you make is overwritten during upgrade and as a result, it is recommended that you back up your dashboards before an upgrade.

This chapter includes the following topics:

■   The Getting Started Page

- Availability Dashboards

- Capacity Dashboards

- Configuration Dashboards

- Cost Dashboards

- Performance Dashboards

- Dashboard Library

- Software Defined Wide Area Network Dashboard

- vRealize Automation 7.x Dashboards

- vRealize Automation 8.x Dashboards

- Service Discovery Dashboards

- Inventory Dashboards

- Microsoft Azure Dashboards

- AWS Dashboards

- Dashboards in VMware Cloud on AWS

- Dashboards in NSX-T Management Pack

The Getting Started Page

Operations Management is a set of interdependent disciplines. Knowing the relationship between these disciplines is as important as knowing each of them separately. The relationship between the disciplines matters because the symptom displayed and the root cause are often two different things, for example, sometimes a configuration problem can lead to a performance problem.

Availability

- Availability considers HA (high availability) settings. As a result, planned downtime (for example, ESXi on maintenance mode) impacts availability.

- Availability, done right, does not impact Capacity and Performance as it is already accounted for.

- The higher the Availability SLA, the higher the price. There is a significant difference for each additional 9 of availability. Five 9s costs a lot more than four 9s.

Performance and Capacity

- Performance is more time sensitive and important than capacity. You must manage performance first and then manage capacity.

- Performance and capacity have an opposite relationship. Highest performance is achieved at lowest capacity, as that is when the VM or the infrastructure is delivering the most amount of work.

- Capacity management is about maximizing utilization, without compromising any performance. It also considers latent workload and future demand.

Cost and Price

- Cost goes hand in hand with capacity. The higher the utilization of the IaaS, the lower the cost per VM. Cost is separate from capacity as it can be optimized without reducing capacity.

- Price can move independent of cost. It has concepts such as discount and progressive pricing. Use price to discourage large unused VMs.

- The better the performance SLA, the higher the price the customer is willing to pay, hence the term Price/Performance.

Compliance and Security

- Compliance is measured against both internal and industry standards.

- Security is related, but not the same as configuration.

Configuration and Inventory

- Inventory is related, but not identical to configuration. Configuration impacts performance, cost, capacity, and compliance. Therefore, it is the primary focus of optimization assessment. Inventory is what you have. Configuration includes properties of what you have. For example, the number of VMs in a cluster are a part of the inventory and not a part of configuration. The number of ESXi hosts in a cluster are a part of inventory and configuration because that is how the cluster is designed. The cluster is configured with eight ESXi hosts for the same reason.

There are two types of counters that impact performance and capacity. Contention is the primary counter for performance, and utilization is the primary counter for capacity. Utilization serves performance and capacity differently. For performance, look at the actual and real utilization. For capacity, it is measured against usable capacity (after HA and a buffer). While they have a negative correlation, contention can develop at low utilization. Unbalanced and configurations are two typical causes of low utilization. Allocation complements demand as newly provisioned VMs tend to be idle (which can last for months). Future load cannot be detected by the demand model as they do not exist. The allocation model should be used to complement the demand model.

The Seven Pillars of Operations Management and the Management Process

The best practice of operations management requires you to distinguish between the pillar and process. The pillar is what you must manage, and the process is how you manage them.

Management Dashboards — The Seven Pillars of Operations Management.	Day 0: Planning — Set your target threshold according to your expectations.	Day 2: Monitoring — Compare the reality with the plan.	Day 2: Troubleshooting — Identify possible issues and resolve them.	Day 2: Optimizing — Reduce cost, increase efficiency, and automate process.
Availability	Yes	Yes	Yes	May be
Performance	Yes	Yes	Yes	Yes
Compliance	Yes	Yes	No	Yes
Capacity	Yes	Yes	No	Yes
Cost	Yes	Yes	No	Yes
Configuration	Yes	Yes	No	Yes
Inventory	No	May be	No	No

When to Manage

What to Manage

Each pillar is an individual unit of management, namely capacity management, performance management, and compliance management. They represent individual disciplines and are compatible with one another. Each pillar's complexity depends on the technology, for example, vSAN's capacity is more dynamic than the central array. In vSAN, changing the storage policy can create a sudden spike.

Day 0 provides the expected result. Some companies conduct a stress test, load test, so they know what to expect when the real load comes in. Without proper planning, you cannot know what the reality is, as you have not defined the process well.

Troubleshooting is an activity and not something you manage. It focuses on the reason, and then formulates a solution to prevent future incidents. Incidents either mean something dead, slow, or breached. You troubleshoot availability, performance, and security.

Inventory is something you have, not something you plan. You plan for capacity, with a certain configuration. Inventory merely accounts for what you have. Nothing to troubleshoot nor optimize.

Using the Getting Started Page

The Getting Started page breaks tasks into broad three broad categories, Management, Flows, and Collections. Use the Getting Started Dashboard to understand the relationship between these categories.

The Management category includes the seven pillars of operations, Availability, Performance, Compliance, Capacity, Cost, Configuration, and Inventory.

The Flows category of dashboards covers the process that includes Troubleshooting, and Optimization. You can use the Troubleshooting dashboards to resolve any potential issues related to availability, contention, utilization, and configuration. Troubleshooting is more than simply identifying the problem. It focuses on the reason behind the problem and also formulates a solution to prevent reoccurrence. An incident means that something is either dead, slow, or has been breached. You can troubleshoot availability, performance, and capacity. Use the Optimization dashboards to enhance the performance of your environment. You can choose to

correct a problem area, update, simplify, or improve your virtual machines and infrastructure. You can optimize performance, capacity, cost, and configuration. You even improve the availability of your system to an extent but you cannot enhance the compliance or inventory.

Lower Cost	• Reclamation: Orphaned VMs, powered off VMs, idle VMs, and oversized VMs snapshots. • Reduce DC Footprint: Save software (MS, Red Hat, VMW), hardware (server, storage, network), and data center (rack, space, cooling, UPS). • Move burst capacity from own to on-demand.
Better Performance	• Performance Profiling: Enable proactive monitoring via actual baseline. • Establish performance SLA that complements availability SLA. • NOC Dashboards: Insights followed by alerts. • Faster business service using self-service and approval workflows.
Lower Complexity	• Standardize architecture. • Standard operating procedure. • Reduced human error due to automation. • Upgrade outdated software and replace ageing hardware.
Higher Customer Satisfaction	• Internal IT department: Reputed among Apps team. • External SP: Repeat business. • Price/Performance: Ability to justify or defend pricing.
Higher Compliance	• Internal compliant (for example, vSphere Hardening). • Industry regulation (for example, PCI DSS, HIPAA).

The Collection category comprises of Public Cloud and the Library sections. The AWS and Azure dashboards are displayed under the Public Cloud dashboards. You can choose to view the overall performance of these services or view specific dashboards related to the services. The Library contains dashboards related to the Network Operating Center and the Executive. It also lists dashboards that do not fit into the pillars of operation, like the VOA and the deprecated dashboards.

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 type and then select a dashboard from the Getting Started page or click the dashboard name listed on the right side of the Getting Started page.

Note Deprecated dashboards are no longer part of the Getting Started page. They can be accessed from the dashboards drop-down menu under Dashboard Library.

Availability Dashboards

Availability covers the uptime of the object now and the uptime trend over time. The availability of hybrid clouds should be tracked at both the provider and consumer layers to understand the availability of the environment. These dashboards show the current uptime and the uptime percentage over the past month.

VM Availability Dashboard

Use the **VM Availability** dashboard to calculate the availability of the Guest OS. The availability of the Guest OS is calculated because the Guest OS might not be running even when the VM is powered on. There are two layers of Availability, that is, the Consumer layer and the Provider layer. This dashboard covers the Consumer layer. You can view VMs in the selected data center, uptime trend for a selected cluster, and so on.

Design Considerations

The **VM Availability** dashboard helps you check the availability (uptime in percentage) of VMs, as availability is typically part of the services provided by the IaaS provider.

This dashboard does not check the application uptime because it is possible that the application such as, a database, or a web server, is down while the underlying Windows or Linux is up. Generally, the service provided by the IaaS team is only for Windows or Linux. For information on the application, use the network ping or application-specific agent such as application monitoring.

How to Use the Dashboard

- In the **Datacenters** widget, click any data center from the list.

  - To view the overall information, click the **vSphere World** object.

  - The other widgets are automatically updated once you click any datacenter.

- Create a filter that reflects your class of service for this widget. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads and you can focus on the important VMs. You can achieve this by creating a vRealize Operations Manager custom group for each class of service.

- The **VMs by Uptime in the last 30 days** widget displays the average uptime of VMs grouped by their availability. The bucket distribution helps you cater to a wide array of environments. If you are monitoring only production VMs where the uptime is expected to be near 100% all the time, edit the bucket to meet your operational needs.

  - The VMs in the **Selected Datacenter** widget display all the VMs that are currently deployed to the data center. The average uptime is displayed for the last month. For a production VM, expect this number to be 100% or closer to 100%.

    Note The Services column will be blank unless Service Discovery is enabled and the services/processes are discovered on a specific virtual machine.

  - The VMs column includes all VMs including the powered off VMs.

- Click any VM in the **VMs by Uptime in the last 30 days** widget to view the details in the **VM in the Selected VM Powered On Status**, **Selected VM Uptime Trend**, and **Selected Cluster Uptime Trend** widgets.

  - The **Selected VM Uptime Trend** widget displays the selected VM's Guest Tool Uptime (%) across the last 30 days.

- The **Guest OS: Services** widget displays the service state over time and the process or services running inside the Guest OS. If Guest OS services or processes are discovered inside a VM, their availability is analyzed. This requires the Service Discovery.

- The **ESXi Host(s) where the VM has run** widget displays the historical migration of the VM. This can be useful in determining the cause of a VM downtime.

Points to Note

- The metric only tracks the availability of VMware Tools and not the entire Guest OS. If VMware Tools is not up, it assumes the Guest OS to be down. You can check that this is not a false negative by adding a few line charts that display the evidence of activity. A good counter is IO counters such as Disk IOPS, Disk Throughput, and Network Transmit Throughput, because IO requires CPU processing. CPU usage is not a reliable counter as the work by VMkernel on the VM is charged to the CPU counters.

- vRealize Operations Manager exhibits a new ping adapter. This allows you to enhance the accuracy of the uptime measurement by creating a super metric that adds the ping information or by checking the process using an agent, such as application monitoring.

- Add a property widget that lists the selected VM properties to give you more context about the VM. In a large environment, the VM name alone might not provide enough context.

vSphere Availability Dashboard

There are two layers of Availability, that is, the Consumer layer and the Provider layer. The **vSphere Availability** dashboard covers the Provider layer. This dashboard includes a cluster and not an ESXi host because the cluster is operationally a single compute provider. This dashboard considers the N+1 design, where the cluster can withstand one host failure. Logically, a cluster with fewer hosts has a higher risk.

Design Considerations

The **vSphere Availability** dashboard helps you analyze and report the uptime, as availability is typically part of the official business SLA. It is also often required in the monthly operational summary report.

This dashboard is not designed for live monitoring of the uptime. A NOC style of dashboard is better suited for those use cases. VMware Tools such as vRealize Log Insight must be leveraged as the fault is typically preceded with soft errors.

How to Use the Dashboard

- The **Clusters** widget lists all the clusters in the environment. It is sorted by the lowest uptime so that the cluster with the lowest uptime in the last one month is displayed.

    - The **Running Hosts** column is color-coded as logically a smaller cluster has a higher risk. A single host failure results in a relatively higher capacity degradation.

    - The **vSAN?** column is hyper-converged, which means both the compute and the storage part is considered.

    - The **Admission Control Policy** column is based on the Cluster Configuration \ DAS Configuration \ Active property. The mapping between the code to name is:

        - -1 : Disabled

        - 0 : Cluster Resource Percentage

        - 1 : Slot Policy (Powered-on VMs)

        - 2 : Dedicated Failover Hosts

    - In a large environment, creating a filter for the list of clusters can make it more manageable. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, you can easily view your gold clusters.

- Click any cluster from the **Clusters** widget.

    - The cluster uptime is automatically plotted in the **Selected Cluster Uptime Trend** widget. It uses 99%, 99.%, and 99.99% as the threshold for red, orange, and yellow colors respectively.

    - The ESXi host details in **ESXi in the Selected Cluster** widget are automatically updated. For more context, you can add a property widget that lists the selected ESXi host properties.

- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.

- The **Datastores not available** widget lists only the datastores with powered off status. This covers both local and shared datastores. To add context, consider adding an extra column such as the data center where it resides, and the datastore types such as NFS and VMFS.

- The **Port Group Availability** widget lists port groups that currently have an uptime of less than 100%. To add context, consider adding an extra column such as the data center where it resides, number of used ports, and the maximum number of ports.

- For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.

- In a large environment, you can create a filter for this dashboard. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads.

- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.

Points to Note

- You can add vCenter Server and NSX components availability. This requires the VMware SDDC Health Monitoring Solution.

Ping Overview Dashboard

Use the Ping Overview dashboard to configure the ping functionality and verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN. You can view ping adapter details like, latency distribution and packet loss distribution in this dashboard.

Customizations Available for Your Use

For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.

Widget Information

- Latency distribution - You can use this widget to see the objects that are experiencing high latency.

- Packet Loss Distribution – You can use this widget to see the objects that are experiencing high packet loss.

- Ping Targets – You can use this widget to view the list of ping targets grouped by their FQDN. Latency and packet loss information is also displayed for the ping objects.

- Breakdown by Source Initiator – You can use this widget to view the List of ping statistics by the source (ping initiator). You can ping the target from multiple locations, to determine if the issue is network-related or server-related.

Capacity Dashboards

Capacity quantifies the resources used, resources remaining, and opportunities to reclaim unused resources. Projections of the demand provide a proactive view of capacity. The **Capacity Dashboards** display capacity in terms of time remaining before capacity is projected to run out, the amount of capacity remaining, the number of VMs that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

Capacity management is about balancing demand and supply. It is about meeting demand with the lowest possible cost.

For IaaS or DaaS, capacity management begins before the hardware is deployed. It begins with a business plan that defines what class of service will be provided. Each class of service, for example, gold, silver, bronze is differentiated by the quality of service and covers the availability, for example, 99.99% uptime for Gold, 99.95% uptime for Silver. It also covers performance, for example, 10 ms disk latency for Gold, 20 ms disk latency for Silver, and security or compliance.

The quality incurs cost and in turn drives price. Gold VM is higher per vCPU and per GB of RAM because it has a higher quality of service. A proper pricing model must be planned. If you want your customers to rightsize in advance, then a 64 vCPU VM has to be more than 64x the price of 1 vCPU VM. If the pricing model is a simple straight line, there is no incentive to go small and no penalty if it is over provisioned. In this case, you end up forcing rightsizing in production, which is a costly and time consuming process.

Demand is more than the active load that is consuming your capacity. Since capacity based on utilization is incomplete by itself, the principles displayed in the following figure are considered.

- Latent demand. Many critical VMs are protected with Disaster Recovery. During a Disaster Recovery drill or an actual disaster, this load is consumed.

- Potential demand. Many newly provisioned VMs take time to reach their expected demand. It takes time for the database to reach the full size, the user base to reach the target, and the functionalities to complete. When this is achieved, it results in the increase in demand.

- Unmet demand occurs when the VM or Kubernetes Pod is undersized. The load is running nearly 100% most of the time.

- Excessive demand can wreak havoc in a shared environment. A group of highly demanding VM can collectively impact overall performance of the cluster or datastore.

Cluster Capacity Dashboard

The **Cluster Capacity** dashboard helps you visualize information differently by providing choices for customization. Use this dashboard to highlight the clusters that need attention. The **Cluster Capacity** dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology.

Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management. The **Cluster Capacity** dashboard considers the following factors that impact capacity:

- Contention

- Utilization

- Allocation

- Reclamation

Contention is included as it directly measures the performance. If your cluster is unable to serve its existing workload, then do not add a new workload. By definition, if the cluster does not have room for a new workload, then its capacity is full. The ideal scenario is that the cluster must run at 100% utilization but 0% contention. In this case, the cluster is productive and your investment is well used.

Utilization is the primary counter for capacity, as it reflects the actual live usage of the resources. When utilization is high, it does not matter if the overcommit ratio is far below your target because the cluster is full. Also, the utilization must not be very low.

Allocation complements utilization as not all workload is real. Some demands can suddenly emerge such as:

- Newly provisioned VM

- Disaster Recovery

- Undersized VM

- Auto-scale VM (a group of web servers behind a Load Balancer)

Reclamation is included as it can impact your decision and the wastage can be common. Capacity can be low, but if you can reclaim a sizeable chunk of wastage, you can defer the purchase of the hardware.

Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **Cluster Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

- The **Clusters by Capacity Remaining** and **Clusters by Time Remaining (days)** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.

- The two bar charts work together. The ideal situation is low capacity remaining and high time remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

- The three heat maps are **Time Remaining**, **Capacity Remaining**, and **VM Remaining**.

- The cluster size is made constant for ease of use. If your cluster sizes are not standardized, consider using the number of ESXi hosts to display the difference in sizes.

The third layer displays a table, accompanied by other widgets to display details of the selected cluster.

- **Clusters Capacity List** widget. If any cluster needs attention, then select the cluster to view the related details.

- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.

- Reservation can impact the efficiency of your cluster. If your cluster size varies, complement the reservation number by displaying a relative value.

- Number of VMs is displayed because the newly provisioned VMs might not be active yet. They are often mistaken as idle, as they can remain unused for months. When you view VM increasing but the demand remaining low, it is a sign of potential demand coming up in the future.

- Workload can be low, but is the overcommit ratio high? The newly provisioned VMs tend to be idle for weeks, and suddenly increase. Use the **VM Count** widget to view if there was recent growth.

- You can check why it is low on capacity. Is it because of real workload or just reservation?

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.

- If you have screen real estate, add a cluster size information. Add cluster size. Small clusters are less efficient from a capacity perspective due to higher overheads and the inability to support larger VMs.

- The peak is defined as the highest among any ESXi hosts. If the peak is higher than the cluster-wide average, then it is unbalanced and is a common reason for suboptimal capacity. You can add a peak to complement the average utilization. Find out the cause of unbalance and optimize it.

- Add peak to complement average utilization. This lets you focus on unbalance, a common reason for suboptimal capacity. Find out the source of unbalance, which can be an opportunity for optimization.

- This dashboard is not designed for the stretched cluster as it requires its own capacity model.

Datastore Capacity Dashboard

The **Datastore Capacity** dashboard highlights the datastores that need attention. This dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology. The **Datastore Capacity** dashboard is designed for both the

VMware administrator and the Storage administrator to foster closer collaboration between the two teams.

Design Considerations

See Capacity Dashboards for common design consideration among all the dashboards for capacity management.

Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

■ The **Shared Datastores by Capacity Remaining** and **Shared Datastores by Time Remaining** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.

■ The two bar charts work together. The ideal situation is low capacity remaining and high time remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

■ There are three heat maps, the primary heat map being **Remaining Capacity**.

■ The other two heat maps **Used Capacity by Datacenter** and **Used Capacity by Datastore Cluster**, cover used capacity. The **Used Capacity by Datastore Cluster** heat map is designed for datastore clusters.

The **Shared Datastores** widget is grouped by data center.

■ If you use Datastore Cluster as your standard, replace the grouping with it. This widget is sorted by the least capacity remaining.

Select a datastore from the **Shared Datastores** widget. The remaining widgets will automatically display the capacity details of the selected datastore.

■ The **Disk Space** widget displays the total capacity allocated and the actual capacity used. You can compare the total Capacity vs the Provisioned capacity vs the Used capacity. If the allocated space increases and the actual capacity does not, then it means that the VMs have not been used. You can ensure that you do not run out of space sooner than expected.

■ In the **VM Count** widget, a rising number that is not complemented by a similar rise in the used space indicates a latent demand.

■ There are three reclamation opportunities: powered off VM, snapshot, and orphaned VMDK.

  ■ The snapshot must be 0 GB. If it is not 0, then it should be temporary. A snapshot lasting beyond a day must be investigated.

- Orphaned VMDK are the ones that are not associated with any VM. The orphaned VMDK must be 0.

Points to Note

Storage in VMware IaaS is presented as a datastore. In a large environment, group datastores as datastore clusters for ease of operations. vSAN uses datastores to present its storage, but it requires a different formula for capacity and performance management. In certain situations, RDM (Raw Device Mapping) and network file shares are also used by certain VMs.

ESXi Capacity Dashboard

The **ESXi Capacity** dashboard supports the **Cluster Capacity** dashboard and is also required for the non-clustered ESXi.

Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

How to Use the Dashboard

The **Summary** heat map provides an overall view of the ESXi Host capacity, grouped by their clusters.

- Each ESXi host is represented by a box, displaying their capacity remaining.

- The ESXi host size is made constant for ease of use. If your ESXi sizes are not standardized, consider using the number of physical cores or Total CPU GHz to display the difference in sizes. Ensure that the smallest ESXi is not too small.

- Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

The **ESXi Hosts Capacity** widget lists all the ESXi hosts in your environment, grouped by their parent cluster.

- The standalone ESXi are displayed at the bottom under No Group.

- In a large environment with many data centers, you can zoom into a specific vCenter or data center. You can also filter or search for specific ESXi hosts matching certain names.

- The **99th Percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric. To rule out the outlier, the worst performance (which is equivalent to 100th percentile) is not considered. Also, the performance threshold is set to be stringent.

Select one of the ESXi hosts from the **ESXi Hosts Capacity** widget. All the three line charts automatically display the trend of selected ESXi host.

- Displays both total and usable utilization in terms of RAM and CPU.

- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.

- A technology refresh is often used to address the capacity shortage. Consider adding a property widget that displays the hardware model and specification to help you determine the age of the hardware.

VM Capacity Dashboard

The **VM Capacity** dashboard provides a quick overview of all the VMs in a given data center and their capacity and time remaining. This dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology.

Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

The **VM Capacity** dashboard helps you visualize information differently by providing choices for customization. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

Select a data center from the **Datacenters** widget.

- The **VMs by Capacity Remaining** bar chart displays the distribution of VMs by capacity remaining in the selected data center. It provides a quick overview of the VMs that are undersized or oversized.

- The **VMs by Capacity Remaining** heat map provides details by grouping the VMs by clusters, so that you can view the clusters that need attention.

- The VM size has been standardized for better visualization. You can add the size that suits your capacity team better.

- The **VMs Capacity in the Selected DataCenter** widget is sorted by the VM with the least capacity remaining. You can sort it by Time Remaining to suit your capacity team better. This table is color-coded.

Select a VM from the **VMs Capacity in the Selected DataCenter** widget. All the remaining widgets automatically display the capacity information of the selected VM.

- The **Disk** widget displays capacity at the Guest OS partition level. There is no overall capacity at VM level because different partitions have different capacity.

Points to Note

■ Use the custom property and add more context to the VM, such as, owner name, clusters where the VM is running, and datastores where the VM files are stored.

VM Reclamation Dashboard

The **VM Reclamation** dashboard helps you manage various types of reclamation that can be done on virtual machines. This dashboard is designed for the Capacity and the Operations team. The reclamation is grouped by buckets. Use this dashboard to view the trend charts that help you analyze the growth over time without changing the context.

Design Considerations

The **VM Reclamation** dashboard helps you visualize information differently by providing choices for customization. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

This dashboard is divided into two sections:

■ The first section provides a summary, giving you an overall picture of reclamation.

■ The second section provides details, giving you the actual VM name to reclaim.

Review the summary information.

■ The summary is presented in the three bar charts, **Count of VM by Snapshot Size**, **Count of Powered Off VM by Disk Space**, and **Count of Idle VM by Memory Footprint**. Each bar chart corresponds to an area you can reclaim.

   ■ The **VM Snapshots** widget corresponds to VMs that are more than a few days old.

   ■ The **Powered Off VM** widget assumes they are back up and it is safe to delete them.

   ■ The **Idle VM** widget helps you reclaim memory but not CPU. The idle VM memory occupies ESXi host's physical memory.

■ Idle VM does not display any CPU usage as there is nothing to reclaim and no CPU is being used. Since the CPU is idle, the only benefit is the overcommit ratio.

■ Memory reclamation is based on the memory footprint at the parent ESXi host. The value inside the Guest is not what is being reclaimed, and so it is irrelevant.

■ Adjust the bucket size to suit your operational requirements.

■ Select any of the VMs in the above widgets to view its trend over time. The trend chart is placed on the same page, so you can review it without opening a new screen. This helps you quickly toggle between VMs.

■ If the snapshot is expanding rapidly, ensure that the VM disk is large (relative to the underlying datastore) as it can fill up the datastore.

Points to Note

There are five areas of reclamation, so start with the easiest first.
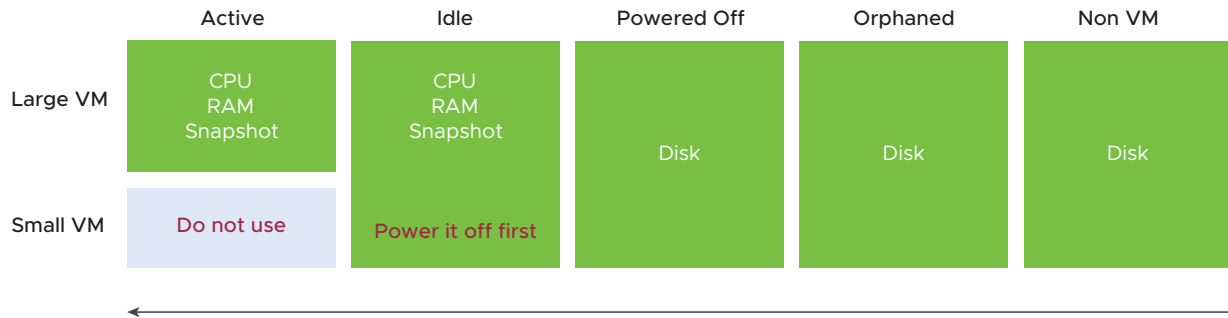
- Non-VM indicates files are not associated with VM. Typically, these are ISO files.

- Orphaned file is a file in the datastore that is no longer associated with any VM. For orphaned Raw Device Mapping (RDM), look from the storage array if there is any ESXi host mounting it. The orphaned VMDK is not listed in this dashboard as it is not associated with a VM. If your environment has orphaned VMDK, add a fourth column in this dashboard.

- Snapshots are not back up and they cause performance problem to the VM. Keep them only for protection during change. Once the change is validated as successful, keeping the snapshot does a disservice to the VM.

- If your environment is large, change the dashboard filter to a functional filter. Group by the class of services such as gold, silver, and bronze and default the selection to the least critical environment. In this way, you can be active in reclamation.

- If reclaiming is a long drawn manual process in your organization, add a filter by department or VM owners. One way to do this is to create a vRealize Operations Manager custom group.

- If the VM name in your environment does not provide sufficient business context, add more information in the table to give context to the VM. Information such as VM owner, clusters where the VM is running, and datastores where the VM files are stored can be useful in the analysis.

- Disk cannot be reclaimed immediately. They have to be in the powered-off stage at least for a week.

vSAN Capacity Dashboard

The **vSAN Capacity** dashboard complements the vSphere **Cluster Capacity** dashboard by displaying capacity related to vSAN. To manage vSAN capacity, use both dashboards.

Design Considerations

As this dashboard is designed to complement the vSphere **Cluster Capacity** dashboard, it shares the same design consideration. It focuses on the storage and vSAN specific metrics but does not list non-vSAN clusters.

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

Wastage is displayed by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **vSAN Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

- The **Clusters by Capacity Remaining** and **Clusters by Time Remaining (days)** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.

- The two bar charts work together. The ideal situation is low Capacity Remaining and high Time Remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

- The three heat maps are **Time Remaining**, **Capacity Remaining**, and **VM Remaining**.

- The cluster size is made constant for ease of use. If your cluster sizes are not standardized, consider using the number of ESXi hosts to display the difference in sizes.

The third layer displays a table, accompanied by other widgets to display details of the selected cluster.

- **vSAN Clusters** widget. If any cluster needs attention, then select the cluster to view the related details.

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.

vSAN Stretched Clusters

The vSAN Stretched Clusters dashboard provides an overview of the cluster resources used across vSAN fault domains. Using the stretched clusters dashboard you can monitor the resource consumption at the site level for Preferred Sites and Secondary Sites. You can create custom dashboards for specific vSAN stretched cluster metrics.

Where to View vSAN Stretched Cluster Objects

On the menu, click **Dashboard > Capacity and Utilization > vSAN Stretched Clusters**.

You can also view the vSAN stretched cluster objects from **Environment > VMware vSAN > vSAN and Storage Devices > vSAN Clusters**, if the vSAN cluster is a stretched cluster.

The vSAN Stretched Clusters dashboard provides information about CPU Capacity, Cores, Memory Capacity, and Disk Capacity for the Preferred Site and the Secondary Site. You can identify the vSAN stretched clusters running out of capacity looking at the utilization metrics.

Configuration Dashboards

As an operations management software, vRealize Operations Manager focuses on the impact to day-to-day operations a product has, rather than the feature of the product itself. Products under monitoring, such as vSphere and vSAN, can have features that are related, but have a different impact on operations. For example, vSphere provides Limits, Reservation, and Shares for the VM.

Limits, Shares, and Reservation. As a feature, they are closely related, and appear in the same dialog box and must be learned as one. However, they impact operations differently. The following table describes that in more detail.

VM Limit	Impacts the VM	• Should not be used. Right-size instead. • Results in unpredictable performance of the Guest OS.	Incorrect Configuration
VM Reservation	Impacts the Infrastructure	• Keep the total amount low, and relative to the total capacity of the cluster. • Absolute value. A 2-GHz reservation is in fact a 2-GHz reservation. • Results in suboptimal infrastructure capacity, as overcommit is not possible.	Suboptimal Configuration
VM Share	Impacts the Infrastructure	• Keep the number of variations to below three. One for each class of service. • Relative value. 2000 worth of reservations depends on the value of other VM reservations. Be careful when you move the VM to another cluster, as the relative value changes. • Results in complex operations. It is harder to troubleshoot performance when the dynamic entitlements of each VM fluctates more.	Complex Configuration

vRealize Operations Manager follows the principle that there are different impacts on operations and applies a methodology for looking at configuration. It does not group the settings by features or objects. Rather, it begins with the impact and prioritizes what can be done.

Correct it?	Update it?	Simplify it?	Optimize it?
Setting is wrong. Correct it.	It is correct, but on an older version. Upgrade or update it.	It is correct and up to date, but it complicates your IaaS operations. Minimize it.	It is correct, up to date, does not complicate your IaaS operations, but it is not as efficient. You can lower cost. Optimize it.
• Settings that are meant to be enabled are disabled. • Insecure settings. • Snapshot running for weeks. • VMs with Limits.	• Running an older version of VMware Tools, Guest OS, ESXi, VMFS, distributed switch, and Telegraf agents. • vSAN: Take the older version.	• Large configuration. Example is a monster VM. • Clusters with too many permutations of shares or VM settings, that override cluster-wide HA/DRS settings.	• Small clusters. HA overhead is relatively higher. • Clusters with lots of reservations. • Powered off ESXi. • Under-sized or over-sized VMs have suboptimal capacity, not configuration management.

Each operation is unique and as a result, customers run operations differently. What is right for other customers, might not be right for you. Even in the same environment, what is right for a development environment might not be appropriate for a production environment.

The following table lists some of the areas for improvement for the operations in your environment:

Areas of Improvement

	Correct it?	Update it?	Simplify it?	Optimize it?
IaaS Consumer: • Process • Applications • Guest OS • Container • VM	• Java JVM or Database \ memory config too large relative to Guest OS • Guest \ Metric not collecting • Guest \ High TX Broadcast packets • VM \ Tools not installed • VM \ Tools not running • VM \ CPU Limit • VM \ Memory Limit • VM \ Old Snapshot • VM \ On local Datastore	• Guest OS \ Tools • Guest OS \ Windows • Guest OS \ Linux • Guest OS \ Telegraf agent • VM \ Hardware (vmx)	• VM \ Large VM (CPU, RAM, Disk) • VM \ lots of disks, NIC card • VM \ lots of IP address. • VM \ with RDM • VM \ on multiple datastores • VM \ Fault Tolerant • VM \ SRM protected • VM \ Hot Add|Remove \ CPU • VM \ Hot Add|Remove \ RAM	• Java JVM or Database \ memory config too small relative to Guest OS • Guest OS \ no visibility • Container \ smaller than the parent VM • VM \ Tools unmanaged • VM \ bigger than the whole ESXi cores. • VM \ bigger than CPU socket. • VM \ Large Snapshot • VM \ Reservation.

IaaS Provider:
• Telegraf
• ESXi
• Cluster
• Datastore & Cluster
• Switch and Port Group
• Hardware
• NSX
• vSAN
	• ESXi \ vMotion disabled
	• ESXi \ Disconnected from vCenter
	• ESXi \ Maintenance Mode
	• ESXi \ NTP disabled
	• ESXi \ Standalone
	• Cluster \ Admission Control disabled
	• Cluster \ HA disabled
	• Cluster \ HA Failover %
	• Cluster \ DRS disabled
	• Cluster \ DRS manual
	• Cluster Inconsistency
	  • BIOS, ESXi: version
	  • BIOS, ESXi: Power Management
	  • ESXi Storage Path
	  • ESXi Hardware
	• Datastore Cluster inconsistency
	  • Capacity
	  • Performance
	• Datastore \ single path
	• Datastore \ no path. This is unlikely.
	• NSX \ no redundancy for Controller, Manager
		• ARC \ server
		• ARC \ agent
		• ESXi \ hardware
		• ESXi \ vSphere
		• ESXi \ 1 Gb NIC.
		• Server \ not on warranty
		• vCenter \ version
		• Datastore \ VMFS version
		• vSAN \ version
		• Switch \ version
		• NSX \ version
			• ESXi \ Too many variations. No standard
			• Cluster \ Many VM Shares (CPU)
			• Cluster \ Many VM Shares (RAM)
			• Cluster \ Resource Pools
			• Cluster \ Stretched compute + storage
			• Cluster \ 32 nodes or more
			• Cluster \ VM to Host affinity
			• Cluster \ Too many storage paths
			• Datastore \ Shared by >1 cluster WLP uses this
			• Datastore \ Many paths
			• Network \ LBT?
			• Network \ MAC Address change
				• ESXi \ low CPU cores count
				• ESXi \ low RAM size
				• ESXi \ Powered Off
				• ESXi \ HT Disabled
				• ESXi \ 4 socket or higher.
				• Cluster \ small clusters \ host especially for vSAN
				• Cluster \ small clusters \ CPU
				• Cluster \ small clusters \ RAM
				• Cluster \ EVC Mode
				• Cluster \ High Reservation
				• Cluster \ DRS Automation Level
				• Cluster \ DPM disabled
				• vSAN \ All Flash: Dedupe disabled
				• vSAN \ All Flash: Compressed disabled
				• Datastore \ small
				• Datastore \ low VM count
				• Datastore \ no ESXi
				• Distributed Switch \ unused

Design Considerations

The dashboards display configurations that need immediate attention, before displaying the overall configuration. This helps you take measures toward optimizing configuration.

Operations vary among customers, and as a result, it is not possible to design one dashboard to meet every customer's operational needs. A configuration that is important for one customer might not be relevant for another customer. Tailor the dashboard to your unique environment. You can collapse or expand the widgets to allow relevant data to be displayed.

The overall layout is designed to balance ease of use, performance (loading time of the dashboard page), and completeness of configuration check. As a result, not all configuration settings are displayed. Lack of screen real estate is another consideration behind the design.

Cluster Configuration Dashboard

Use the **Cluster Configuration** dashboard to view the overall configuration of vSphere clusters in your environment, especially configurations that need attention.

Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

The **Cluster Configuration** dashboard is organized into sections for ease of use.

- The first section of the dashboard consists of three bar-charts. They correspond to the three main features of vSphere clusters, namely High Availability (HA), Dynamic Resource Scheduler (DRS), and Distributed Power Management (DPM).

    - HA: The best practice is to enable HA admission control. You can specify the admission control policy in the vCenter Server and the threshold for failover shares.

    - DRS: The best practice is to enable DRS. Envision the vSphere cluster as a single logical computer that balances within itself.

    - DPM: The best practice is to enable DPM in an environment where environmental concern is the top priority or the high peak rarely occurs as most of the time you run very low utilization.

- The second section of the dashboard consists of eight pie-charts. They show the relative distribution of key configurations.

    - Two of the bar-charts cover admission control. You must enable admission control. The pie-charts displays the policy code instead of the policy name, as it is based on the property: `Cluster Configuration | Das Configuration | Active Admission Control Policy`. The mapping between the code to name is:

        - -1 = Disabled

        - 0 = Cluster Resource Percentage

        - 1 = Slot Policy (Powered-on VMs)

- - 2 = Dedicated Failover Hosts

  - There are two bar-charts that cover the HA Failover Share. One for CPU, and one for memory.

  - The next two bar-charts cover DRS settings. You might want to fully automate DRS, which means that there is no operator intervention required for both initial VM placement and subsequent load balancing, but with a moderate migration threshold (value = 3.0). The value ranges from 1.0 to 5.0.

  - There are two pie-charts that show reservation. One for CPU and one for memory. Minimize the total reservation value as it prevents overcommit of resources and hence results in less optimal utilization. Memory reservation can remain and occupy the memory space of the ESXi host, even though the VM does not use the memory anymore. Consider the analogy of unused files that you have not opened for months in the c:\ drive of your laptop. They still take up space on the hard disk. Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service.

- The third section of the dashboard consists of two bar-charts. They show the absolute distribution of the clusters.

  - The first bar-chart displays the cluster grouped by the number of ESXi hosts. Small clusters, defined as having a lower number of ESXi hosts, have a higher overhead while large clusters have a higher risk if there are cluster-wide outages. Performance risk is lower, because there are more nodes that DRS can tap on, but if there is an actual problem, troubleshooting can be tougher, because there are more nodes to analyze. For large clusters, have a disaster recovery plan as an unexpected cluster-wide outage can impact many VMs.

- The fourth section of the dashboard lets you drill-down to an individual cluster.

  - A table lists all the clusters with their key configuration. You can export this list as a spreadsheet for further analysis or reporting.

  - Select a cluster. The list of ESXi hosts under the cluster, with shares and resource pool information, is automatically filled up.

  - Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service. Avoid providing different services to individual VMs as that increases the complexity of the cluster performance.

  - Keep the number of resource pools minimal.

  - Some of the columns are color coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.

- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.

- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

ESXi Configuration Dashboard

Use the **ESXi Configuration** dashboard to view the overall configuration of the ESXi hosts in your environment, especially the configurations that need attention.

Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

The **ESXi Configuration** dashboard is organized into sections for ease of use.

- The upper section of the dashboard displays basic ESXi configurations that should be standardized for ease of operations.

  - There are six pie-charts that are displayed as one set because there is a relationship between their values. There should be a correlation between them. Ideally, the ESXi version, the ESXi build, and the BIOS must be identical across all ESXi hosts in a cluster. Keep the variations of the hardware model, NIC speed, and storage path minimal. The more complex the pie-chart, the more variants you have. This results in complex operations, that potentially results in higher operating expenses.

  - The configurations should reflect your current architecture standard. Each pie-chart counts the occurrences of a particular value. A large slice signifies that the value is the most common value, and if that is not your current standard, then you must address it.

- The second section of the dashboard displays configurations that are potentially suboptimal.

  - The three bar-charts display various size dimensions of the ESXi hosts. The bar-charts are designed to be seen as one set. Ensure that there are a minimal number of variations to reduce complexity.

- Smaller ESXi hosts have a relatively higher overhead, and are limited in running larger VMs. If they have a low core count, they might be using an outdated CPU. Small ESXi hosts are more expensive on a per core, per GB, per rack unit basis than larger ones if they occupy the same space. However, a 4-CPU socket ESXi host is likely to be too large, resulting in a concentration risk (too many VMs in a single ESXi host). Maintain a good balance that balances your budget and risk constraints.

  - Adjust the distribution chart bucket size to fit your environment.

- The third section of the dashboard displays configurations that you might want to avoid.

  - The six bar-charts focus on security, availability, and capacity settings that you can set as a standard. For example, you should consider enabling the NTP daemon for a consistent time, which is critical for logging and troubleshooting.

  - The three tables list the actual ESXi hosts that are in a non-productive state. They can be on maintenance mode, powered off, or in a disconnected state.

- The last section of the dashboard displays all the ESXi hosts in your environment.

  - You can sort the columns and export the results into a spreadsheet for further analysis.

  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result, there is nothing to display.

- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.

- For complete visibility, consider adding physical server monitoring by using the appropriate management pack. For more information, see the following page.

Network Configuration Dashboard

Use the **Network Configuration** dashboard to view the overall configuration of vSphere distributed switches in your environment, especially for the areas that need your attention.

Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

How to Use the Dashboard

The dashboard is organized into two sections for ease of use.

- The first section displays network configurations that need your attention.

    - There are five bar-charts that focus on critical security settings.

    - The last bar-chart displays the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.

- The second section provides overall configuration information, with the ability to drill down to a specific switch.

    - Click the row to select a switch from the list.

    - The ESXi hosts, port groups, and the VMs on the switch are displayed.

    - Review each of the tables. For the ESXi host table, ensure that the settings are consistent.

    - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

    - You can sort the columns and export the result into a spreadsheet for further analysis.

Points to Note

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.

- For complete visibility, consider adding physical network device monitoring by using the appropriate management pack. For more information, see the following page.

- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

VM Configuration Dashboard

Use the **VM Configuration** dashboard to view the overall configuration of virtual machines in your environment, especially for the areas that need attention.

Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

- Click the row to select a data center from the data center table.

  - In a large environment, loading thousands of VMs increases the web page loading time. As a result, the VM is grouped by data center. In addition, it might make sense to review the VM configuration per data center.

  - For a small environment, vSphere World is provided, so you can view all the VMs in the environment.

  The **VM Configuration** dashboard is organized into three sections for ease of use. All the three sections display the VM configuration for the selected data center.

- The first section covers limits, shares, and reservations.

  - Their values can easily become inconsistent among VMs, especially in an environment with multiple vCenter Servers.

  - Shares should be mapped to a service level, to provide a larger proportion of shared resources to those VMs who pay more. This means that you should only have as many shares as your service levels. If your IaaS provides gold, silver, and bronze, then you should have only three types of shares.

  - Value of the shares and reservation is relative. If you move a VM from one cluster to another (in the same or different vCenter Server), you might have to adjust the shares.

  - Reservation impacts your capacity. Memory reservation works differently from CPU reservation, and it is more permanent.

- The second section covers VMware Tools.

  - VMware Tools is a key component of any VM, and should be kept running and up to date.

- The third section covers other key VM configurations.

  - Keep the configurations consistent by minimizing the variants. This helps to reduce complexity.

  - **VM Network Cards** widget. If you suspect that your environment might have a VM with no NIC, consider adding it as a dedicated bucket.

- The last section of the dashboard is collapsed by default.

  - You can view all the VMs with their key configurations.

  - You can sort the columns and export the results into a spreadsheet for further analysis.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.

- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

- The pie-chart and bar-chart cannot drive other widgets. For example, you cannot select one of the pie-slices or buckets, and expect it to act as a filter to a list or a table.

- You can apply a specific color in a pie-chart or distribution chart for a specific numeric value, but not string value. For example, you cannot apply the color red to the value `Not Installed`.

vSAN Configuration Dashboard

The **vSAN Configuration** dashboard provides overall configuration details and is useful in large clusters with many vSANs, where you have to follow a certain standard configuration.

Design Considerations

See Configuration Dashboards for common design considerations among all the dashboards for configuration management.

How to Use the Dashboard

The **vSAN Configuration** dashboard is organized into three sections for ease of use.

- The first section displays six pie-charts.

  - There are five bar-charts that focus on critical security settings.

  - The last bar-chart shows the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.

- The second section displays three bar-charts.

  - The three bar-charts together provide a good overview of the vSAN key capacity configuration. By analyzing the distribution, you can identify if you have capacity configuration that is outside your expectation.

- The last section of the dashboard displays all the vSAN clusters with their key configuration.

  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

  - You can sort the columns and export the result into a spreadsheet for further analysis.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

Workload Management Configuration Dashboard

This dashboard provides a quick configuration summary of all the key objects associated with workload management such as Supervisor Clusters, Namespaces, vSphere Pods and Tanzu Kubernetes clusters. It is essential that the configuration is consistent across all the objects. Configuration drifts may result in inconsistent performance or availability of the applications leveraging workload management Kubernetes constructs.

Use the dashboard to ensure that the configuration is consistent across all objects.

You can view the following widgets in the dashboard.

- **Environment Summary**

- **Supervisor Cluster Versions**

- **Cluster Status**

- **Pod Data**

- **Supervisor Cluster Configuration Summary**

- **Pod Configuration Summary**

- **Kubernetes cluster Configuration Summary**

- **Namespace Configuration Summary**

Consumer \ Correct it? Dashboard

The **Consumer \ Correct it?** dashboard complements the main VM configuration dashboards by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The dashboard is designed to focus on VMs that need attention. Lists are used to keep it simple, and show actual objects. The lists can be tailored using the filter and the custom group. The lists can also be exported for an offline discussion.

The dashboard is extendable, reflecting the reality that different customers have a different set of settings to verify. Since the dashboard layout is a collection of tables (List View), you can extend it by adding more tables. You can add more List View widgets to verify the VM configurations that your operations require.

How to Use the Dashboard

The **Consumer\Correct it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Tools Widgets:

    - Using VMware Tools has multiple benefits. For the list of benefits, refer to KB 340.

    - vRealize Operations Manager uses VMware Tools to retrieve Guest OS metrics. Without this, right-sizing VM memory can be inaccurate, because the hypervisor metrics (VM Memory Consumed and VM Memory Active) are not designed to measure Windows or Linux memory utilization. ESXi VMkernel does not have visibility into the Guest OS for security reasons.

    - Independent software vendor (ISV) support is the most common reason that VMware Tools is not installed. The ISV vendor might claim that no additional software is installed in their appliance unless they have certified it. For more information about VMware Tools, see the VMware Tools documentation.

    - If VMware Tools is installed, there might be reasons why the application team disables it. The Infrastructure team should inform and educate their application team, and document the technical recommendations about why VMware Tools is recommended to be running all the time.

- CPU Limits and Memory Widgets:

    - It is recommended that you do not use memory and CPU limits as it can result in an unpredictable performance. The Guest OS is not aware of this restriction as it is at the hypervisor level. It is recommended that you shrink the VM instead.

- Guest OS Counters Missing Widget:

    - There is no visibility into the Guest OS performance counters because the requirements are not met. The memory counter is especially important as VM Consumed and VM Active are not replacements for Guest OS counters. See KB 55675 for more details.

- Old Snapshot Widget:

    - Ensure that the snapshot is removed within one day after the change request. If not, it might result in a large snapshot and impact the performance of the VM.

Points to Note

- Add a banner summary to the top of this dashboard so that you can verify if there is an incorrect confirmation. Add a scoreboard and select the World object and then collapse all the tables below. Create a super metric for each summary and apply it to the World object.

- In a large environment, create a filter for this dashboard to enable you to focus on a segment of the environment. Group it by a class of service such as, gold, silver, and bronze. Default the selection to gold, your most important environment. In this way, your monitoring is not cluttered with less critical workloads.

- There are other VM configurations that maybe relevant to your environment. Review the list of VM settings that you might want to add to this dashboard.

- For context, add a property widget that lists the selected VM properties. In this way, you can check the property of your interest without leaving the screen. Multiple List View widgets can drive the same property widget, so you do not have to create one property widget for each List View.

- If your operations require it, add a list of VMs that do not have these three key performance counters: CPU Run Queue, CPU Context Switch, and Disk Queue Length.

Consumer \ Optimize it? Dashboard

The **Consumer \ Optimize it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Optimize it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities. A suboptimal configuration might not impact performance or increase complexity, but it can be more expensive.

Design Considerations

The **Consumer \ Optimize it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How To Use the Dashboard

The **Consumer \ Optimize it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- VM Reservation:

  - VM reservation causes a positive impact on the VM, but a negative impact on the cluster. Total reservation cannot exceed cluster capacity. This creates a suboptimal cluster as VMs do not use the entire assigned memory at the same time.

  - VM reservation places a constraint on the DRS placement and HA calculation. Avoid using reservation as a means to differentiate performance SLA among all the VMs in the same cluster. It is difficult to correlate CPU Ready with CPU Reservation. A VM CPU Ready does not improve two times because you increase its CPU reservation by two times. There is no direct correlation.

- Guest OS visibility:

    - Since your workloads are sharing resources and are over-committed, your operations are easier if you know what is running inside. This helps with monitoring and troubleshooting, resulting in optimal operations.

    - For critical VMs, consider logging the Guest OS, such as Windows and Linux, to capture errors that do not surface as metrics. These errors typically appear as events in the log files or in the event database in the case of Windows. Use vRealize Log Insight to parse Windows events into log entries that can be analyzed.

- Snapshot:

    - Old snapshots tend to be larger. They consume more space and have a higher chance of impacting performance.

Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result shares limitations and customization ideas.

Consumer \ Simplify it?

The **Consumer \ Simplify it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Simplify it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Consumer \ Simplify it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The **Consumer \ Simplify it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Large VMs (CPU, Memory, and Disk):

    - A large VM, relative to the underlying ESXi host and datastore, requires more careful planning (Day 0) and monitoring (Day 2).

- Ensure that the VM size does not exceed the size of the underlying ESXi host. If your ESXi host has CPU hyper-threading, do not count the logical processor. Instead, count the physical core. For best performance, keep it within a (non-uniform memory access) NUMA boundary.

- During monitoring, verify if the VM is highly utilized. If the VM vCPU count is equal to the ESXi cores, and the VM is running at almost full capacity, you might not be able to run other VMs. Large VMs can impact the performance of other VMs, especially if it is given higher shares. Only when the large VM is under-utilized, can the ESXi hosts run other VMs.

- If the number of configured vCPUs on a VM is higher than the number of cores per socket on the ESXi, the VM can experience the NUMA effect. If the ESXi has more than one physical CPU (socket), cross-NUMA access negatively impacts performance.

- The larger the VM, the longer the time required to vMotion, Storage vMotion, and backup.

- For disk space, if the disk is thin-provisioned and under-utilized, you can deploy other VMs in the same datastore. Ensure that the snapshot is tracked closely, as the risk of capacity running out is higher for a large virtual disk.

- VMs with many virtual disks:

  - It is simpler to have a 1:1 mapping between Guest OS partitions and the underlying virtual disk (VMDK or RDM).

  - For performance and capacity, evaluate the disks and partitions. Each virtual disk must be monitored in terms of IOPS, throughput, and latency. Having multiple virtual disks increases the monitoring and troubleshooting need.

  - If the reason for having many virtual disks is performance, identify which counter serves as proof that multiple virtual disks are required. It is possible that the performance required is met by a single virtual disk.

- VM with many IP addresses or NICs:

  - A VM might need multiple networks, such as production, back up, and management. It is recommended that you route the network interfaces through the NSX-Edge VM. A VM that has multiple network interfaces can bridge the network, causing security risks or network problems.

  - A VM that is part of multiple networks can do so with just a single NIC. A single NIC can be configured to access multiple networks, with each interface having their own IP configuration.

Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

Consumer \ Update it? Dashboard

The **Consumer \ Update it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Update it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Consumer \ Update it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The **Consumer \ Update it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated Tools Widget:

  - Lists all the VMware Tools versions that are still supported. Tailor the filter to fit your operational needs.

- Outdated VM Hardware Widget:

  - Lists all the VM vmx versions that are not 13, 14, 15, or 16. Tailor the filter to fit your operational needs.

- Outdated Windows and Red Hat Widgets:

  - Lists all the Windows client versions that are not version 10.

  - Lists all the Windows server versions that are not versions 2016 and 2019.

  - Lists all the RHEL versions that are not version 7 or 8.

  - If you run other operating systems like Ubuntu, clone the widget. You can also repurpose the widget if you do not run RHEL and Windows.

Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

Provider \ Correct it? Dashboard

The **Provider \ Correct it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard

is designed for vSphere administrators and the platform team. The **Provider \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Provider \ Correct it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations.

  - A cluster is the smallest logical building block for compute. Consider it as a single computer with physically independent components. As a result, consistency matters.

  - Clusters with DRS set to manual. This means that DRS initiated vMotion does not take place unless the administrator manually approves it. Since DRS calculates every five minutes, your quick approval is required to prevent a change of condition.

  - Clusters with HA disabled. Without high availability provided by the infrastructure, each application must protect itself from an infrastructure failure.

  - Clusters with DRS disabled. DRS focuses on performance and capacity, while HA focuses on availability. Without DRS, you must build a buffer on every ESXi host to cope with peak demand.

  - Clusters with Admission Control disabled. Reservation is respected only when Admission Control is enabled.

- The second section covers the ESXi host configurations.

  - ESXi with Network Time Protocol disabled. Logs are a critical component of operations, and are the main source of information in troubleshooting. While troubleshooting performance across objects, the sequence of logs determines which event is the likely root cause, as the oldest event starts the chain of events.

  - A disconnected ESXi host indicates that the ESXi host is not participating in HA and you cannot migrate any VM on it.

  - An ESXi host that is in maintenance mode does not contribute resources to the cluster or the data center if there is a standalone ESXi.

- The third section covers ESXi host configurations that must be consistent within a cluster.

  - BIOS version and ESXi versions.

  - BIOS Power Management, ESXi: Power Management. Ideally, should be set to OS controlled. The ESXi level should be set to balance level.

- ESXi Storage Path. Ensure that the number of paths and the path policies are identical.

- ESXi hardware specifications. Different specifications can result in inconsistent performances experienced by the VM.

Points to Note

- See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

- If you have a standalone ESXi, and you plan to replace it with a clustered ESXi host, add a table to list them.

- Based on your security settings, add a table to check the Distributed Switch and Port Group to ensure that security settings such as promiscuous mode, are used correctly.

Provider \ Optimize it? Dashboard

The **Provider \ Optimize it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Optimize it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

Design Considerations

The **Provider \ Optimize it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations:

  - A small cluster has a higher HA overhead when compared to a large one. For example, a three-node cluster has 33% overhead while a 10-node cluster has 10%. For vSAN, a low number of hosts limits the availability option. Your choice of FTT is relatively more limited.

  - Many small clusters result in silos of resources. As a cluster behaves like a single computer, ensure that it has enough CPU cores, CPU GHz, and Memory. For ESXi in 2020, it is typical to have 512 GB of RAM. This results in 12 TB of RAM for a 12-node cluster, which is enough for DRS to place many VMs as it balances them.

  - If there is a lot of reservation, add a list for clusters with a relatively high reservation. If your clusters are of different sizes, use a super metric to convert the reservation value to a percentage.

- The second section covers ESXi host configurations.

  - Small ESXi. A small host faces scalability limits in running a larger VM. While a 2-socket, 32-cores, 128 GB memory ESXi can run 30 vCPU, 100 GB RAM VMs, the VM experiences a non-uniform memory access (NUMA) effect.

  - ESXi powered off. You can mark the ESXi hosts for decommissioning using the custom property feature of vRealize Operations Manager. You can then create a separate list, so they are not overlooked.

- The third section cover storage and network.

-

- Unused network (distributed port group). This is a potential security risk as you might not monitor it.

Points to Note

- See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

- For CPU cores, a change in vSphere licensing means that the ideal core is 32-cores per CPU socket. This maximizes the software license. For more information, see the vSphere Pricing Model.

Provider \ Simplify it? Dashboard

The **Provider \ Simplify it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Simplify it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

Design Considerations

The **Provider \ Simplify it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

- Click the row in the **Clusters** widget to select one of the clusters from the table.

  - A cluster is more complex to operate when it has resource pools, shares, and limits.

- Review the list of resource pools:

  - Ensure that the number of VMs in each resource pool reflects the intended settings for the VM. The resource pool value is divided and shared among the VMs. The more the VMs, the lesser the resources allotted to each VM.

- Verify if there are VMs who are siblings to the resource pools.

  - Verify if the resource pools are further split into subresource pools.

- Review the CPU Share and Memory Shares pie-charts:

  - Multiple combinations of shares, especially both CPU and memory, makes troubleshooting difficult.

  - Each share must map to exactly one class of service, such as one for gold and one for silver as the shares define the class of service. Shares are also relative, meaning the value depends on the value of sibling objects, such as, resource pool or VM. Ensure that the values are consistent across clusters to avoid unintended consequences while moving the VM to another cluster.

- Review the CPU Reservation and Memory Reservation tables:

  - High total reservation, especially both CPU and memory, complicates the cluster operations as it impacts the HA slot calculation, and limits the DRS choice of placement.

- Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

Provider \ Update it? Dashboard

The **Provider \ Update it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Update it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

As part of operations best practices, keep the infrastructure up to date. Running outdated components that are too far behind the latest version, can cause support problems or upgrade problems. It is common that the fix for the problem is only available in the later versions. Outdated hardware can also result in higher operating costs. Outdated hardware might cost more data center footprint, such as rack space, cooling, and UPS. Refreshing your technology and consolidation are two common techniques to optimize cost.

Design Considerations

The **Provider \ Update it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboards

The **Consumer \ Update it?** dashboard is a collection of tables (List View) that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated vSphere Components Widgets:

  - Lists all the vCenter Servers versions that are not 6.7 or 7.0.

  - Lists all the ESXi host versions that are not 6.5, 6.7, or 7.0.

  - Lists all the vSAN ESXi host versions that are not 6.7 or 7.0. A more stringent filter is applied for vSAN because of a relatively higher maturity in the latest release. From vRealize Operations Manager and vRealize Log Insight, there are more counters, properties, and events that improve monitoring and troubleshooting.

  - Lists all the vSphere distributed switches, regardless of the version.

  - You should tailor the filter to fit your operational needs.

- Outdated Server BIOS Widget:

  - Lists all the ESXi hosts regardless of the BIOS version. Edit the widget and tailor the filter to fit your operational needs.

- Other than customizing the existing widgets, consider adding the following checks:

  - ESXi hosts with outdated hardware, using a filter based on your environment.

  - ESXi hosts that are no longer on warranty. Create a custom property to capture the end of warranty.

  - Physical storage arrays with outdated firmware, model, and an expiring warranty.

  - Physical network switch with an outdated OS version and hardware model

    Note Install the relevant management pack for the last two points.

Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

Cost Dashboards

The dashboards in the cost category cater to cloud administrators who are responsible for managing the expenses related to your cloud infrastructure. Using Cost dashboards, you can compare the cost of VMware cloud infrastructure with other cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently.

Assess Cost Dashboards

The **Assess Cost** dashboard provides an overview of the scale of your infrastructure in terms of physical capacity available.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- You can view the total cost of ownership per month for the infrastructure and the savings opportunities details, if any, for the infrastructure.

- You can view the details of the division of infrastructure investments across all data centers. The dashboard provides the magnitude of each data center in terms of the number of physical servers and virtual machines. It also provides details about the amount of savings that can be achieved from each of these data centers.

- The dashboard displays data about how you invest across clusters of different quality offered across all vCenter Servers.

Base Rate Analysis Dashboard

The **Base Rate Analysis** dashboard helps you analyze the cost efficiency of your data center.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- The total cost of ownership is the cost required to run your data center per month. This is derived from the cost drivers.

- The average cost per VM is derived by considering the cost of all the VMs in your environment. The cost of each VM depends on the base rate of the cluster the VM is placed on and its utilization. The base rate of the cluster is computed based on the total cost of ownership and the expected utilization levels of the cluster. Storage base rates are directly obtained from cost drivers.

- If the cluster is running on an allocation-based capacity model, the base rate is derived from the total cost of the cluster and the over-commit ratio. The base rate is indicative of how costly a resource is, on a given cluster.

- A base rate is derived from the total cost and the expected utilization of the cluster.

- A deeper analysis of the base rates can be performed using the CPU, memory, or storage-related widgets, which help rank clusters and datastores relative to their base rates.

Datacenter Cost Drivers Dashboard

The **Datacenter Cost Drivers** dashboard provides the cost of different data centers in a private cloud.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- You can select individual data centers to view summary and trends. The summary of the data center costs is grouped into two:

  - Compute. Covers all the costs that are spent on compute related hardware, software, and services.

  - Non-Compute. Covers storage and network.

- Expense trends provide cost variations over a period which indicate infrastructure additions or removal to the data center.

- Cluster expenses indicate the component clusters of a data center that consume the costs. Datastores that represent the storage part of the data center cost are listed alongside.

  Note Network costs are mapped directly to ESXi hosts and hence are costed under compute as well, as of today. This might change in the future.

- When you select a cluster, you can view the component hosts that the cluster is made up of and their monthly depreciated costs. It also provides details on the purchase cost of the server and how many months until it depreciates completely.

Note Server costs can be suggested out-of-the-box by the system, or can be customized by the user. Depreciation information is not available for servers when the server costs are suggested out-of-the-box by the system. Depreciation information is available for those servers when the server cost is customized by the user.

Showback Dashboard

The **Showback** dashboard helps you navigate between multiple groups that you might want to perform showback for. Select an object to view total costs of the group and potential savings within the group.

Customizations Available for Your Use

The definition of a group is based on several constructs such as the vCenter folder, the vRealize Automation 7 business group, and the vRealize Automation 8 project among others. To change this to your definition, edit the widget and select the desired object types under **Output Filter > Basic > Object Types**.

Widget Information

- The **Top Expensive** and **Top Savings** widgets point out the most costly VMs and the VMs with the most savings opportunities within the group.

- A detailed list of the members helps you perform a deeper analysis. You can also export the details into a CSV for offline analysis.

- When you select a particular VM in the list, you can analyze the cost trends for the VM, and as a result, find out the root cause for VM cost variations.

Performance Dashboards

Performance is about ensuring workloads get the necessary resources. Key Performance Indicators (KPI) can be used to identify performance problems related to workloads. Use these KPIs to define SLAs associated with tiers of service. These dashboards use KPIs to display the performance of workloads at the consumer layer and the aggregate performance of workloads at the provider layer.

SLA is the formal business contract that you have with your customers. Typically, SLA is between the IaaS provider (the infrastructure team) and the IaaS customer (the application team or business unit). Formal SLA needs operational transformation, for example, it requires more than technical changes and you might need to look at the contract, price (not cost), process, and people. KPI covers SLA metrics and additional metrics that provide early warning. If you do not have an SLA, then start with Internal KPI. You must understand and profile the actual performance of your IaaS. Use the default settings in vRealize Operations Manager if you do not have your own threshold, as those thresholds have been selected to support proactive operations.

The following graphics depict the above relationship.

Formal SLA

SLA is part of the business contract. Customers can track their SLA using the self-service portal.

One policy for each SLA.

Internal KPI

IaaS Performance is Quantified and Measured.

Performance is based on the Production Environment.

Default Policy.

Reactive

Complaint-based Operations.

Blamestorming.

IaaS Performance Measures based on Business Impact.

The Three Processes of Performance Management

In performance management, there are three distinct processes.

- Planning. Set your performance goals. When you architect a vSAN, you must know how many milliseconds of disk latency you want. 10 milliseconds measured at the VM level (not the vSAN level) is a good start.

- Monitoring. Compare the plan with the actual. Does the reality match what your architecture was supposed to deliver? If not, you must fix it.

- Troubleshooting. When the reality is not according to the plan, you must fix it proactively and not wait for issues and complaints.

To understand what is not healthy for performance management consider the following areas in the given order.

1   Contention: This is the primary indicator.

2   Configuration: Check the version incompatibilities.

3   Availability: Check for soft errors. vMotion stun time, lock up. This requires Log Insight.

4   Utilization: Check this in the end. If the first three parameters are good, you can skip this.

The Three Layers of Performance Management

There are three main realms of enterprise applications. Each of these realms has its own set of teams. Each team has a set of unique responsibilities and requires the associated skill set. The three realms comprise of Business, Application, and IaaS. Refer to the graphic below to understand the three layers and the typical questions asked on each layer.

Layers		Sample Metrics		
Business	Business Result	• How many sales did we make today? • How many customers bought our product this week? • On an average, how long did the XUYZ transaction take in this hour? • How many customers logged in yesterday? • On an average how long did customers stay logged in?		
	Business Transaction			
Application	Individual Node	• How long did the SQL query ABCD take in the last 7 days ? • One hour ago, what was the value of the SQL server free memory ? • What is the overall application uptime? • Are my applications configured for performance?	Vertical Metrics depend on each application and its needs	2
	The System			
IaaS	VM or Container	• What is the Windows CPU Run Queue? • In the past 24 hours, what was the peak VM CPU contention? • What was the total number of IO hitting vSAN from 9am - 6pm yesterday? • What is the buffer in a physical switch right now?	Horizontal Common metrics are applicable for all applications	1
	Virtual Infra			
	Physical Infra			

Performance Management is largely an exercise in elimination. The methodology slices each layer and determines if that layer is causing the performance problem. Hence it is imperative to have a single metric to indicate if a particular layer is performing or not. This primary metric is aptly named Key Performance Indicator (KPI).

The upper layer depends on the layer below it, and hence the infrastructure layer is typically the source of contention. As a result, focus on the bottom layer first, as it serves as the foundation for the layer above it. The good part is this layer is typically a horizontal layer, providing a set of generic infrastructure services, regardless of what business applications are running on it.

The Two Metrics of Performance Management

The primary counter for performance is contention. Most look at utilization, because they fear something wrong might happen if utilization is high. That something is contention. Contention manifests in different forms like, queue, latency, dropped, canceled, and context switch.

However, do not confuse ultra-high utilization indicators as a performance problem. If your ESXi host experiences ballooning, compression, and swapping, it does not mean that your VM has a performance problem. You measure the performance of the host by how well it serves its VMs. While performance is related to the ESXi host utilization, the performance metric is not based on the utilization, instead it is based on contention metrics.
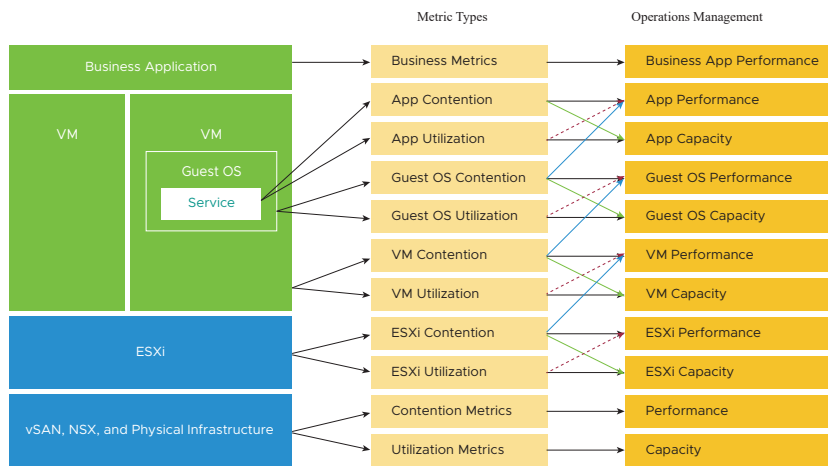
It is possible for VMs in the cluster get affected from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at an individual consumer (VM). The following table shows various possible reasons.

Infra Configuration	VM and Guest OS Configuration
ESXi Settings ■ Host and BIOS power management causes Frequency to drop. ■ HT enabled. It looks like twice the capacity, but it is actually 1.25 X throughput. ■ ESXi - HW compatibility. Driver and firmware are two areas that can impact the performance. ■ Mismatch of queue depths along the various storage stacks. Must calibrate all the way to the physical array. ■ vMotion too slow or high stunned time.	VM: Limit, Share, and Reservation ■ Make sure that no limit is set. CPU ready includes limit. ■ Make sure that the shares are consistent (as per what the VMs want or you agree to.) ■ Avoid reservation if possible. This impacts the net available resources for the other VMs.
Network ■ MTU mismatch. ■ Hops. Especially horse-shoe, or going through multiple ESXi.	Size: NUMA effect. VM spanning NUMA nodes.
Cluster Settings ■ Inconsistent configuration among hosts in a cluster. EVC Mode can play a part if the hosts are from different generations. ■ Resource Pool 　■ Make sure the shares match the number of VMs. 　■ Make sure that no VM is siblings to RP. ■ VM- Host Affinity. ■ DRS Setting.	Snapshot. IO is processes 2x. VM drivers.
vSAN ■ The host where the storage was having performance issues.	Windows or Linux process ping pong, process runaway, and OS level queue.

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

VM Performance

Since VM is the most important object in vSphere, it warrants an extra explanation. The graphic below lists the counters you should look at.

	CPU	RAM	Network	Disk
Inside Guest OS (Linux, Windows) Need VMware Tools	Run Queue Context Switch	Paging Rate (MB/s) Committed %	OS Output Queue Length Driver Queue	OS Queue Driver Queue
	Utilization	In Use Modified + Standby	Throughput (Mbps) Latency	Latency
Outside Guest OS (Guest OS can't control)	Run I Used System + VMX +MKS	Active, Consumed, Granted, Swapped-in	Throughput	IOPS, Throughput (Large Block)
	Ready + Co-Stop + Overlap IO Wait + Swap Wait	Contention	TX Dropped Packet Normalized Latency	Outstanding IO Latency

The KPI counters can get technical for some users, so vRealize Operations include a starting line to get them started. You can adjust the threshold, once you profile your environment. This profiling is a good exercise, as most customers do not have a baseline. The profiling requires an

	Metric	Green	Yellow	Orange	Red
Guest OS Contention	Total CPU Run Queue	0-5	> 5	> 10	> 20
	CPU Context Switch Rate	0 - 5K	< 25K	< 100K	> 100K
	Total Disk Queue Length	0 - 25	> 25	> 50	> 100
Guest OS Usage	RAM Free (MB)	> 512 MB	> 256	> 128	≤ 128
	RAM Page-in Rate (KB/s)	0 - 25K	> 25 K	> 50 K	> 100K
VM Contention	CPU Co-Stop (%)	0 - 2.5%	> 1	> 3	> 5
	[SLA] CPU Ready (%)	0 - 2.5%	> 2.5	> 5	> 7.5
	Total CPU Overlap (ms) at VM level	0 - 1000	> 1000	> 2500	> 5000
	CPU IO Wait	0 - 1000	> 1000	> 2500	> 5000
	[SLA] RAM Contention (%)	0 - 1%	> 1	> 2	> 4
	[SLA] Disk Latency (ms)	0 - 10 ms	> 10	> 20	> 40
	[SLA] Network TX Dropped Packet	0	> 0	> 1	> 2
VM Usage	CPU Usage (%)	0 - 85%	> 85	> 90	> 95

advanced edition.

Performance Metrics

vRealize Operations Manager uses the following threshold for internal KPI.

IaaS	VM Counter	Threshold
CPU	Ready	2.5%
RAM	Contention	1%
Disk	Latency	10 ms
Network	TX Dropped Packet	0

The table is an example of a stringent threshold. A high standard for performance is used because it is an internal KPI for the consumption of the infrastructure team. It is not an external formal SLA that is confirmed with the customers. There must be a buffer between the internal KPI and the external SLA so that the operations team receive early warnings and has the time to react before the external SLA is breached. A high standard also works from the mission critical point to view to the development environment. If the standard is set to the least performing environment, then it cannot be applied to the more critical development.

A single threshold is used to keep the operations simple. This means that the performance in production is expected to have a higher score than the development environment. The development environment performance is expected to be worse than the production environment, while everything else is equal. A single threshold helps to explain the difference in Quality of Service (QoS) provided by a different class of service. For example, if you pay less, you get a poor performance and if you pay half the price, expect to get half the performance.

The four elements of IaaS (CPU, RAM, Disk, and Network) as mentioned in the table, are evaluated on every collection cycle. The collection time is set at five minutes as it is an appropriate balance for monitoring. If SLA is based on one minute, it is too close and results in either cost increase or reduction in threshold.

Design Considerations

All the performance dashboards share the same design principles. They are intentionally designed to be similar, as it is confusing if each dashboard looks different from one another, considering they have the same objective.

The dashboards are designed with separate two sections: summary and detail.

- The summary section is typically placed at the top of the dashboard to provide the overall picture.

- The detail section is placed below the summary section. It lets you drill down into a specific object. For example, you can get the detailed performance report of any specific VM.

In the detail section, use the quick context switch to check the performance of multiple objects during performance troubleshooting. For example, if you are looking at the VM performance, you can view the VM-specific information and the KPIs without changing screens. You can move from one VM to another and view the details without opening multiple windows.

The dashboard uses progressive disclosure to minimize information overload and ensure the webpage loads fast. Also, if your browser session remains, the interface remembers your last selections.

Many of the performance and capacity dashboards share a similar layout since there is a shared commonality between these pillars of operations.

Guest OS Performance Profiling Dashboard

Use the **Guest OS Performance Profiling** dashboard to know the actual performance of your environment.

Some counters directly impact the performance of Windows or Linux, the operating systems running inside the VM. These KPIs are outside the control of the hypervisor.

Modern operating systems such as Linux and Windows use memory as cache, since it is faster than a disk. Some counters directly impact the performance of Windows or Linux. These KPIs are outside the control of a hypervisor, which means that the ESXi VMkernel cannot control the increase or decrease of the KPI values. The KPI visibility also requires an agent, such as VMware Tools. As a result, they are typically excluded in performance monitoring.

Since they are closer to the applications, it is critical to know their values and establish an acceptable range. The acceptable level of these KPIs among all the VMs in your environment varies. By profiling the actual performance across time and from all VMs, you can establish a threshold that is supported by facts. Since there are 8766 instances of 5 minutes in a month, profiling 1000 VM over a month means you are analyzing 8.8 million datapoints.

Design Considerations

The dashboard uses progressive disclosure to minimize information overload and ensures that the webpage loads fast.

In a large environment, loading thousands of VMs increases the loading time of vRealize Operations Manager. As a result, the VM is grouped by data center. For a small environment, vSphere World is provided so you can see all the VMs in the environment.

How to Use the Dashboard

Select data center from the data centers list. The three tables listing CPU, memory, and disk will show the VMs in the selected data center or vSphere world. Each table shows the highest value in the last one week (2016 datapoints based on five minutes collection cycles), and hence uses the term max as a prefix, for example Max Page-Out/sec or Max Guest OS Disk Queue.

Select any of the VMs in any of the tables. The three line charts are displayed. They are showing data from the same VM to facilitate correlation.

- CPU table widget:

    - The Max CPU Queue column shows the highest number of processes in the queue during the given period. As a best practice, keep the queue below three for each queue. A VM with eight CPUs has eight queues, hence keep this number below 24.

    - The CPU Hyperthreading gives twice the queue as it should as both threads are interspersed in the core pipeline.

    - CPU Context Switch. There is a cost associated with the context switch. There is no guidance for this number, and it varies widely.

- Memory list widget:

    - In memory paging, the modern operating systems (Linux and Windows) use memory as cache, it is much faster than a disk. It proactively pre-fetches pages and anticipates future needs (Windows calls this Superfetch). The rate pages that are being brought in and out can reveal memory performance abnormalities. A sudden change, or one that has sustained over time, can indicate page faults. Page faults indicate that pages are not readily available and must be brought in. If a page fault occurs too frequently, it can impact application performance. While there is no concrete guidance, as it varies by application, you can view a relative size. operating systems typically use 4 KB or 2 MB page sizes.

- Disk list widget:

  - Disk queues are queued IO commands that are not sent to the VM. They have been retained inside the Guest OS (either at a kernel level or a driver level). A high disk queue in the guest OS, accompanied by low IOPS at the VM, can indicate that the IO commands are stuck waiting on processing by Windows/Linux. There is no concrete guidance regarding these IO commands threshold as it varies for different applications. You should view this with the Outstanding Disk IO at the VM layer.

Points to Note

- These Guest OS widgets do not appear unless the vSphere pre-requisites are met. For more information, see KB article 55697.

- Once you determine an acceptable threshold for your environment, consider adding thresholds to the table so you can easily view the VMs that exceed a threshold.

- The CPU queue is the sum from all virtual CPUs. A larger VM can tolerate a higher queue as it has more processors. If you want to compare VMs of different sizes, create a super metric that calculates the queue per vCPU. For more information, see Create a Super Metric.

- Group the VM by clusters of the same class (for example, Gold), so you can see the profile for each environment.

- For a smaller environment, consider changing the table from listing data centers to listing clusters.

Network Top Talkers Dashboard

Use the **Network Top Talkers** dashboard to monitor network demand in your IaaS. In a shared environment, a few VMs generating excessive activity can impact the entire data center. While a single VM might not cause a serious problem, a few of them can.

Design Considerations

The **Network Top Talkers** dashboard helps you analyze how hard these VMs hit your IaaS. It classifies the workload into two: short bursts and sustained hits. A short burst lasts for a short period, maybe for a few minutes. A sustained hit can last for an hour and cause serious problems.

The **Network Top Talker** dashboard forms a pair with the **Storage Heavy Hitter** dashboard. To understand the IO demand in your environment, use both of them concurrently.

The **Network Top Talkers** dashboard displays sustained hits that last for an hour, as they can cause serious problems in a shared IaaS environment. You can identify the villain VMs and compare their demands with the capabilities of the underlying IaaS.

How to Use the Dashboard

The dashboard shows the current workload. This is the total network load (received and transmitted) from all the vSphere environments monitored by vRealize Operations Manager. The idea is to give you an indicator on how hard the overall load is.

- Select a data center from the data centers list.

    - The columns show the number of clusters, ESXi hosts, and VMs for each data center. The VM count includes the powered off VM. To only see the running VM count, edit the widget.

    - If you want to see information from all the data centers, select the vSphere world row.

    - Upon selection, the Total Demand Line chart and the Top Talkers tables fill up.

- Total Demand Line Chart

    - The total throughput (received and transmitted) in the selected data center.

    - Displays both, the five minute peak and the hourly average in one line chart. You can click the metric name to hide it.

- Top Talkers Table

    - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important. For example, an ESXi with 2 x10 GB port can theoretically handle 20 GB TX + 20 GB RX as its full duplex.

Points to Note

- Understanding high demand helps you monitor IaaS and plan your capacity. IaaS provides four services, CPU, memory, disk, and network. While CPU, memory, and disk are bound, an active VM can consume all your network bandwidth, packet per second capacity, and the storage IOPS capcity. A VM with 4 vCPU and 16 GB memory cannot consume more than this amount, the same applies to disk space. A VM configured with 100 GB disk space cannot consume more than that.

- Network throughput, disk throughput, and disk IOPS can spike as their physical limits are very high per VM. This means that IaaS has enough capacity for all workloads and performs well until the VMs start consuming abnormally high amounts of network and disk bandwidth.

Storage Heavy Hitters Dashboard

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard. To understand the IO demands in your environment, use both of them together. If you are using ethernet-based storage, storage traffic runs over the same physical network as your ethernet-based network traffic.

Design Considerations

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard, so they share a consideration behind their design. For more information, see Network Top Talkers Dashboard.

How to Use the Dashboard

- See the **Network Top Talkers** dashboard as they have the same design.

  - The main difference between **Storage Heavy Hitters** and **Network Top Talkers** is that the storage IO has two dimensions: IOPS and throughput.

  - Network IO does not have the IOPS dimension as the packet size is identical (1500 bytes being the standard packet, and 9000 bytes being the jumbo frames).

  - Storage IOPs and throughput are related, so use both to gain insight, they should display a similar pattern. If not, that indicates varying block sizes. For example, a throughput spike without an accompanying IOPs spike indicates large block sizes.

- Which VMs hit the storage the hardest.

  - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important, because different classes of SSD have different IOPS and throughput capabilities.

    After identifying the villain VM, talk to the VM owners if the numbers are excessive during peak hours and identify the reasons behind the excessive usage. You must ensure that they do not create a hot spot. For example, vSAN cluster with > 100 disk can handle numerous IOPS but if the VM objects are only on a few disks, those disks can become a hot spot.

Points to Note

- Interpreting IOPs and throughput metrics depends on your underlying physical storage. For visibility into this hardware layer, add physical storage metrics to the dashboard.

VM Contention Dashboard

The **VM Contention** dashboard is the primary dashboard for VM performance. It is designed for VMware administrators or architects. It can be used for both, monitoring, and troubleshooting. Once you determine that there is a performance issue, use the **VM Utilization** dashboard to see if the contention is caused by high utilization.

Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

For understanding the performance concept of the selected counters and their thresholds, see the Performance Dashboards

How to Use the Dashboard

- Select a data center from the data center table.

  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

    Note   The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- The two bar charts are automatically shown.

  - Use them together to get an insight about your CPU readiness and your Memory contention analysis. Analyze how the cluster serves the VMs. For each VM, it picks the worst metric in the last 24 hours. By default, vRealize Operations Manager collects data every 5 minutes, so this is the highest value among 288 datapoints. Once it has the value from each VM, the bar charts puts each VM in the respective performance buckets. The threshold in the buckets considers best practices, hence they are color coded.

  - For any critical environment, expect that all the VMs are served well by the IaaS. You must see green on both distribution charts. For development purposes, you can tolerate a small amount of contention in both CPU and Memory.

- VM Performance in selected Data Center.

  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.

  - The table is sorted by KPI Breach columns, directing your attention to the VMs that are not served well by the IaaS.

  - The table shows the hostnames known by Windows or Linux. This is the name that the application team or VM owner knows, as they might not be familiar with the VM name.

  - The rest of the columns show performance counters. Because the goal is proactive monitoring, the counters are the worst and not the average, during the monitoring period. Because the operations context here is performance, not capacity, the table considers the last 24 hours only. Daily use is encouraged as any activity older than 24 hours is considered irrelevant from a performance troubleshooting viewpoint.

  - The column KPI Breach counts the number of SLA breaches in any given 5 minutes. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter

varies from 0–4, with 0 being the ideal. The value 4 indicates that all 4 IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, not an external SLA. Your internal threshold should be more stringent, so that you have a reaction time.

- Select a VM from the table.

    - All the health charts show the KPI of that VM.

    - The health charts display the last value, lowest value, and the peak value. Expect that the peak is within your threshold.

Points to Note

- This dashboard uses Guest OS counters and VM counters appropriately. The two layers are distinct layers, and they each provide a unique visibility that the other layers might not give. For example, when the VMkernel de-schedules a VM as it has to process something else (for example, other VM, kernel interrupt). The Guest OS does not know the reason. In fact, it experiences frozen time for that particular vCPU running on the physical core and experiences time jumps when it is scheduled again.

- Guest OS counters logically require VMware Tools.

- The health chart is color coded. Change the settings if it does not suit your environment. If you are unsure of what suitable numbers to set for your environment, profile the metrics. The Guest OS Performance Profiling Dashboard dashboard provides an example of how to profile metrics.

- For a smaller environment with one or two data centers, change the filter from data center to cluster. Once you are list a cluster, you can then add the cluster performance (%) metric and sort them in an ascending order. This way the cluster that needs immediate attention is on the top.

- If you have a screen real estate, group the VMs by cluster or by ESXi host. This way, you can quickly see if the problem is in a particular cluster or ESXi host.

- Change the default timeline from one week to one day as and when required to suit your operations.

- If you navigate a lot to the **VM Utilization** dashboard from this dashboard, add a connection using the dashboard to dashboard navigation feature. For more details, see Dashboard Navigation Details.

VM Utilization Dashboard

The VMware administrator uses the **VM Utilization** dashboard with the **VM Contention** dashboard for managing performance.

Design Considerations

Use the **VM Utilization** dashboard to identify virtual machines with a high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when a queue develops inside the Windows or Linux operating systems. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- Select a data center from the data center table.

  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

    Note The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- VM Peak CPU Usage (%).

  - There is no peak memory use as it is not applicable. Memory is a form of storage, for example consider a hard disk occupied space. A 90% utilization of the total space is not slower than 10%. This means that the issue is related to capacity issue and not performance.

  - The bar chart is color coded using five colors instead of four. The color gray is introduced to convey any wastage. Resources that are hardly utilized do not signify that the performance is at its peak. It can also mean the opposite. For example, if a VM needs 1+ vCPU, configuring it with 2 CPUs results in better performance instead of configuring it with 128 CPUs.

- VM Peak Utilization.

  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.

  - The table focuses on peak utilization, because the context is performance and not capacity.

- Select a VM from the table.

  - All the health charts show the KPI of that VM.

  - Compliment the free memory with the memory IOPS or the memory throughput metric. The metrics in a gigabyte measure the space, and not the speed. Memory is a form of storage, so what you must measure is the rate, for example, read-write per second.

Points to Note

- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the VM Contention Dashboard.

Troubleshoot an Application Dashboard

The VMware vRealize Application Management Pack provides discovered applications to be managed in vRealize Operations Manager. Using the **Troubleshoot an Application** dashboard, users can see the applications and the relevant metrics and alerts for the selected application. The dashboard also displays its relationship to the infrastructure. In the list of metrics, select a metric to see its trend over time.

Cluster Contention Dashboard

The **Cluster Contention** dashboard is the primary dashboard for vSphere cluster performance. It is designed for VMware administrators or architects. It can be used for both, monitoring and troubleshooting. Once you determine that there is a performance issue, use the **Cluster Utilization** dashboard to see if the contention is caused by high utilization.

Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

Utilization of the cluster is not shown in the **Cluster Contention** dashboard. You must separate the two concepts: utilization and contention. Performance and capacity are different concepts managed by two separate teams. Both CPU and memory are also shown separately. You can have a problem with one, without any issue in the other. CPU is more common as memory tends to have a lower overcommit ratio.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- Average Cluster Performance (%).

    - This is the primary KPI for your entire IaaS. It plots how your IaaS is performing every 5 minutes, giving you the trend view of the overall performance.

    - The metric itself is simply the average of the Cluster KPI / Performance (%) metric. This performance metric in turn averages the VM Performance / Number of KPIs Breached metric from all the running VMs in the cluster. Hence a value of 100% indicates that every running VM in the cluster is served well.

    - As this KPI takes into account every running VM in your environment, the number should be steady. The analogy in real life is the stock market index. While individual stocks can be volatile, overall the index should be relatively steady on a 5 minutes by 5 minutes basis.

- The relative movement of the metric is as important as the absolute value of the metric. Your absolute number might not be as high you want it to be, but if there are no complaints for a long time, then there is no urgent business justification to improve it.

- Clusters Performance.

    - It lists all the clusters, sorted by the least performing cluster in the last one week. You can change this time period.

    - The worst performance shows the lowest number in the time period. As vRealize Operations Manager collects data every 5 minutes, there are 12 x 24 x 7 = 2016 data points in a week. This column shows the worst point among these 2016 datapoints.

    - A single number among 2016 datapoints can be an outlier that needs to be complemented with another number sometimes. A logical choice is the average of these numbers. For the average performance to be low, a lot of criterias have to be low. Waiting for the average causes a delay in your operations, and rise in complaints. For performance monitoring, the 95th percentile is a better summary than the average.

    - Your cluster should function at a 100% and perform its fuctions as planned.

- Select a cluster from the table.

    - All the health charts show the KPI of the selected cluster.

    - For performance, it is important to show both the depth and breadth of the performance problems. A problem that impacts one or two VMs requires a different troubleshooting than a problem that impacts all the VMs in the cluster.

    - The depth is shown by reporting the worst among any VM counter. So the highest value of VM CPU Ready, VM Memory contention, and VM Disk Latency among all the running VMs are shown. If the worst number is good, then you do not need to look at the rest of the VMs.

    - A large cluster with thousands of VMs can have a single VM experiencing poor performance while 99.9% of the VM population is fine. The depth counter might not report that most VMs are fine. It only reports the worst. This is where the breadth counters come in.

    - The breadth counters report the percentage of the VM population that is experiencing performance problem. The threshold is set to be stringent, as the goal is to provide early warning and enable proactive operations.

Points to Note

It is possible for VMs in the cluster to suffer from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at individual consumer (VM). The following table shows various possible reasons.

Event	Aware?
Power Management	No
HT	No
Ready	No
Co-Stop	No
System	No
Steal	No
IO Wait	No
Memory Wait	No

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

Certain settings such as DRS automation level and the presence of many resource pools can impact performance. Consider adding a property widget to show the relevant property of a selected cluster, and a relationship widget to show resource pools.

For a large environment with many clusters, add a grouping to make the list more manageable. Group it by class of service, so you can focus more on the critical clusters.

Cluster Utilization Dashboard

The VMware administrator uses the **Cluster Utilization** dashboard with the **Cluster Contention** dashboard for performance management.

Design Considerations

This dashboard supports the **Cluster Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when VMs experience a contention. By default, vRealize Operations Manager has a 5-minutes collection interval. For five minutes, there may be 300 seconds worth of data points. If a spike is experienced for a few seconds, it may not be visible if the remaining of the 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- CPU(%) and Memory (%).

    - Review the CPU and Memory distribution charts for an overview of the CPU and memory utilization of the clusters.

    - The highest metric in the last one week is used. Average or 95th percentile is not used as this is utilization and not contention. High utilization does not mean bad performance.

    - One week is used instead of one day to give you a longer time horizon and covers the weekend. Adjust the timeline as you deem fit for your operations.

    - Expect memory to be higher than CPU, as it is a form of cache. The Memory Consumed counter is used as it is more appropriate than the Memory Active counter.

    - Low utilization can actually indicate bad performance, as not much of real work gets done. The chart uses the dark gray color for low utilization.

- Clusters Utilization.

  - The cluster utilization table lists all the clusters, sorted by the highest utilization in the last one week. If the table displays the green color, then there is no need to analyze further.

  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.

- Select a cluster from the table.

  - All the utilization charts show the key utilization metrics of the selected cluster.

  - For memory, the high utilization counters are explicitly shown, Balloon, Compressed, and Swapped. Notice they exist even though utilization is not even at 90%, indicating high pressure in the past. If you look only at utilization, you might think you are safe.

  - The line charts show both average and highest among ESXi hosts in the cluster. The reason is unbalanced and it is not rare. There are many settings that can contribute to it (for example, DRS settings, VM Reservation, VM – Host Affinity, Resource Pool, Stretched Cluster, and Large VMs).

  - The disk IOPS is split into read and write to gain insight into the behavior. Some workload is read oriented, while others are write oriented.

  - The disk throughput is not shown as it sums all the traffic. In reality, each ESXi host has its own limit.

  - The vMotion line chart is added, as a high number of vMotion can indicate that the cluster load is volatile, assuming the DRS Automation level is not set to the most sensitive setting.

Points to Note

- If your operations team have some forms of standardization that utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold.

- Consider adding a third distribution chart. Show the balloon counter in this third chart, as it complements the consumed counter. If there is no ballooning, a high consumed value is in fact better than a lower value.

- The workload metric can exceed a 100% because it is demand / usable capacity * 100. This can happen if you have four hosts in a cluster with each host running at 100% demand and admission control is set to 50%.

- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the Cluster Contention Dashboard.

VM Rightsizing Dashboard

The **VM Rightsizing** dashboard helps you adjust the VM size for optimal performance and capacity. It covers both undersized and oversized scenarios. This dashboard is designed for the Capacity and the Operations teams, as rightsizing VMs benefit the day-to-day performance.

Design Considerations

The **VM Rightsizing** dashboard helps you visualize information differently by providing choices for customization. It focuses on a summary that helps discussions with senior management. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

Select a data center from the **Datacenters** widget.

- The cluster capacity remaining is displayed to give a better context. Focus on reclaiming the cluster that is low on capacity remaining, and on upsizing the cluster with a high capacity remaining.

Once you select a data center from the **Datacenters** widget, all the remaining widgets automatically display the information of the selected data center.

- There are two widgets for upsizing recommendation, one for CPU and one for Memory.

- There are two widgets for downsizing recommendation, one for CPU and one for Memory.

- The business processes for oversized and undersized VMs are different, as one requires the affected VM to be shut down and the Owner to return the resources. To upsize, you must add incrementally. To downsize, you must remove in one change window as the effort to reduce is the same and there will be only one downtime.

Points to Note

- The metrics used are `Summary|Oversized|Virtual CPUs` and `Summary|Undersized|Virtual CPUs`. It stores the capacity engine calculation on the recommended number of vCPUs that must be removed or added.

- When you change the VM configuration, the application setting might have to change. This is especially on applications that manage its memory (for example, database and JVM), and schedules the fixed number of threats.

- Avoid reducing vCPUs from greater than 1 to 1 for Windows. The SMP kernel is activated during the first installation and the performance can degrade on a uni-processor machine.

- You can enable Hot Add on VM, but take note of the impact on NUMA.

- For more details on rightsizing, see Rightsizing VMs with vRealize Operations.

Datastore Performance Dashboard

Use the **Datastore Performance** dashboard to view performance problems related to storage such as high latency, high outstanding IO, and low utilization. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- Select a data center from the data centers table.

  - The list of shared datastores in the data center is shown with their KPI.

    Note Datastores that are unavailable are not shown.

- Datastore Performance.

  - Read and Write latency are shown separately in the datastore performance table for a better insight. The nature of read and write problems might not be the same so it is useful to see the difference.

  - Both the worst (peak) performance and 95th percentile are shown. If the latter is close to the peak and high, then it is a sustained problem. If the latter is low, then it is a short duration.

  - The table is color-coded. If your operations require a different threshold, edit the widget to adjust it accordingly.

- Select a datastore you want to troubleshoot.

  - Its read latency, write latency and outstanding IO is automatically shown.

    Note The latency is the normalized average of all VMs in the datastore.

  - Its IOPS and throughput are also displayed. These line charts are not color-coded as it varies per customer. Edit the widget and add your expected threshold. It makes it easier for the Operations team.

  - The list of VMs is displayed.

- Select a VM you want to troubleshoot.

  - Its read latency and write latency are shown.

    Note The number is at the VM level. If you suspect one of the virtual disks has a high latency, use the counter Peak Virtual Disk Read Latency (ms) and Peak Virtual Disk Write Latency (ms).

Points to Note

- The vSphere storage is represented as a datastore. The underlying storage protocol can be files (NFS) or blocks (VMFS). vSAN uses VMFS as its consumption layer as it is unique to vSAN, and has its own monitoring need. Latency can happen when IOPS and throughput are not high. When latency occurs, troubleshooting can take much time.

- You can look at the logs and queue in the various storage stacks (for example, driver) and monitor their performance.

- Datastores that share the same underlying physical array can experience problem at the same time. The underlying array can experience a hot spot on its own, as it is made of independent magnetic disks or SSD.

- The dashboard does not have datastore clusters. If your environment uses it, add a View List to list them, and use this view list to drive the Datastore Performance view list.

ESXi Contention Dashboard

The **ESXi Contention** dashboard is the primary dashboard for managing ESXi host performance. The VMware administrator or architect can use it to monitor and troubleshoot any performance issue. If you determine that there is a performance issue, use the **ESXI Utilization** dashboard to see if the cause for the contention is high utilization.

Design Consideration

The **ESXi Contention** dashboard complements the Cluster Contention Dashboard, and shares the same design consideration.

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed to be used daily, hence the views are set to show data in the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- ESXi CPU Performance and ESXI Memory Performance.

  - Review the two distribution charts for an overview of all the ESXi host's utilization and memory performance.

  - Both charts are using the percentage of VM facing performance counter and not the worst performance among VM counter because you are looking at the ESXi performance and not at the single VM performance. See how it handles all the VMs.

  - The bar chart is color coded. Keep the percentage of the VM population not being served under 10%.

- ESXi Hosts Performance.

  - The ESXi hosts performance table lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is displays the green color, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance greater than 24 hours are likely to be irrelevant.

  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.

- Select an ESXi host from the table.

    - All the health charts show the KPI of the selected cluster.

    - For performance, it is important to show both depth and breadth of a performance problem. A problem that impacts one or two VMs require a different troubleshooting than a problem that impacts all VMs in the cluster.

    - Worst CPU overlap among VMs in the host is included as it indicates a lot of interruptions. A running VM might get interrupted because the VMkernel needs the physical core to run something else. High and frequent numbers of interruptions are not healthy and can impact the VM performance.

    - Expect the network error to be 1% and dropped packet to be 0 most of the times, if not always. If it is not zero, analyze it to see if there are any patterns across all ESXi hosts, and bring it up with your network team.

Points to Note

- Consider adding a third distribution chart and display the CPU co-stop counter in this third chart, as it complements the CPU ready counter. If your environment has relatively slow network and storage IO, you can add IO wait too.

- Unlike the **Cluster Performance** dashboard, there is no average ESXi hosts performance (%) at the vSphere World level. The reason is most ESXi hosts are part of a cluster and monitoring should be done at the cluster level.

- Certain settings such as power management and hyper threading can impact the performance. Consider adding a property widget to show relevant properties of a selected ESXi host.

ESXi Utilization Dashboard

The VMware administrator uses the **ESXi Utilization** dashboard with the **ESXi Contention** dashboard to manage performance.

Design Considerations

The **ESXi Utilization** dashboard supports the **ESXi Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

The dashboard complements the Cluster Utilization Dashboard dashboard, by providing the extra details. Hence it has a similar layout.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- ESXi Hosts Utilization.

  - It lists all the ESXi hosts, sorted by the highest utilization in the last one week. If the table is all displaying the green color, then there is no need to analyze further.

  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.

- Select an ESXi host from the table.

  - All the utilization charts display the key utilization metrics of the selected cluster.

  - For memory, the high utilization counters are explicitly shown, for example balloon, compressed, or swapped. You might notice they exist even though utilization is not even at 90%, indicating that there was a high pressure in the past. If you look at only utilization, you might think you are safe.

  - The disk IOPS and the disk throughput are split into read and write to gain an insight into the behavior. Some workload is read oriented, while others are write oriented.

  - The network throughput is split into sent (transmit) and received to gain insight into the behavior. The total usage can be misleading because it sums up the send and receive traffic. In reality the network pipe is one for each direction (due to the full duplex nature of Ethernet), and not shared.

Points to Note

If your operations team have some forms of standardization that the utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold. For more information, see the points to note in the ESXi Contention Dashboard.

Network Performance Dashboard

Use the **Network Performance** dashboard to view performance problems related to network such as high latency, frequent retransmit, and many dropped packets. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The dashboard enables you to drill down from the distributed switch to the ESXi host and port groups in the switch, and then to the VM.

How to Use the Dashboard

- Distributed Switches.

  - The distributed switches table lists all the switches, sorted by the highest packet dropped. The table splits the incoming traffic and the outgoing traffic for better analysis.

  - As the focus is on performance and not capacity, the throughput counters are not shown.

- Select a switch from the distributed switches table.

  - The health chart shows the dropped packet trend over time.

  - It does not narrow down the list of port groups automatically, as the list of port groups are always showing all the port groups in your environment.

  - If necessary, expand the two collapsed widgets. They show the network throughput and broadcast packets. Utilization is also shown so that you can correlate and understand whether the dropped packets are due to higher utilization.

- Port Groups and ESXi Hosts in the selected switch.

  - They get listed when you select a switch from the distributed switches table.

  - Just like the distributed switch, you can also see their relevant counts.

- If your environment has unused network switches, you can filter them out from this list, as this dashboard focuses only on performance.

Points to Note

- vSphere network is by nature distributed. Each ESXi contributes to the physical NIC. This represents the physical capacity. Distributed switch and its port groups span across these independent network cards. This makes it harder to define and measure its performance. An unbalance can happen among ESXi hosts or physical NIC. In a sense, it is like distributed storage (vSAN). Capacity management does not apply to a port group, since its upper limit (also known as the physical capacity) can vary by even a minute.

- Latency within a data center should be below 1 millisecond. Use vRealize Network Insight to study the latency or the retransmitting problems, caused by moving into the lateral traffic.

- Add a physical network using the appropriate management pack.

Most packets are unicast, between a pair of sender and receiver. If your environment has many VMs sending broadcast packets to everyone and multicast packets to many targets, add a Top-N widget to find out which VMs are sending these packets.

vSAN Contention Dashboard

The **vSAN Contention** dashboard is the primary dashboard for managing vSAN performance. The VMware administrator or architect can use it to monitor and troubleshoot the vSAN cluster performance. If you determine that there is a performance issue, use the **vSAN Utilization** dashboard to see if the cause for the contention is high utilization.

Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The **vSAN Contention** dashboard complements the ESXi Contention Dashboard, and shares the same design consideration. It focuses on the storage and vSAN specific metrics, and does not repeat what is already covered. It does not list any non vSAN cluster.

How to Use the Dashboard

- vSAN Peak VM Latency, vSAN Peak CPU Ready, vSAN Peak Dropped Packet.

  - Review the three distribution charts for an overview of all the vSAN clusters performance.

  - The vSAN peak VM latency chart shows the distribution of disk latency experienced by all the VMs in the cluster. You should expect most of the VMs to experience latency that matches your expectation. For example, in an all flash systems, the VMs should not have >20 ms disk latency. If your vSAN environment is all flash, you must adjust the distribution bucket to a more stringent set.

  - The vSAN peak CPU ready chart shows if any of the vSAN kernel modules has to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not wait for CPU time. vSAN gets higher priority than VM World as it lives in the kernel space.

  - The vSAN peak dropped packet chart shows if any of the vSAN clusters are dropping packet in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

- vSAN Clusters.

  - It lists all the vSAN clusters, sorted by the least performing.

  - It lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is showing all green, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance issues greater than 24 hours are likely to be irrelevant.

  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.

- Select a vSAN cluster from the vSAN clusters table.

  - All the health charts show the KPI of the selected cluster.

  - If you are using SMART, the two heat maps at the bottom of the dashboard provide early warning.

Points to Note

- A large vSAN cluster can have many components. Each of these components can have multiple performance metrics. The total number of KPI can reach hundreds of metrics. For example, take a 10 node cluster. It can have 530 counters to check. vRealize Operations Manager aggregates them by introducing a set of KPIs. This analysis reduces the number to a more manageable number. The following table shows the KPIs and their formula.

Name	What it is
Max Capacity Disk Latency (ms)	Highest latency among all capacity disks take the worst, not average, as the latency in a single capacity disk is already an average of all its VMs. If there are 50 VMs on the disk and 30 are issuing IO on it, then its average is among 30.
Min Disk Group Write Buffer Free (%)	Lowest free capacity among all the disk group write buffers. If this number is low, one of your buffers is not enough. While you want to maximize your cache, a low number is an early warning for capacity management.
Max Disk Group Read Cache/Write Buffer Latency (ms)	Each disk has a Read Cache Read Latency, Read Cache Write Latency (for writing into cache), Write Buffer Write Latency, and Write Buffer Read Latency (for de-staging purpose). This takes the highest among all these four numbers and the highest among all disk groups. It is the max of the max because each of the four datapoints is an average of all the VMs on it.
Sum Disk Group Errors	Sum of the bus reset + sum of commands canceled among all the disk groups. You must use sum and not get the max as each member should return zero.
Count Disk Group Congestion Above 60	The number of disk groups congestion greater than 60. 60 is hardcoded in the vSAN Management Pack as it is a good starting point. As any congestion above 60 serves an early warning, count how many of such occurrences happen.
Max Disk Group Congestion	The highest congestion among all disk groups. A high number indicates that at least one disk group is not performing.
Min Disk Group Capacity Free (%)	The lowest free capacity among all disk groups. A low space triggers rebalance.
Min Disk Group Read Cache Hit Rate (%)	The lowest hit rate among the disk group read cache. Ensure that this number is high as it indicates that the read is served by cache.
Sum vSAN PortGroup Packets Dropped (%)	Sum of all vSAN VMkernel port RX dropped packet + TX dropped packet. You should expect no dropped packet in your vSAN network.

vSAN Utilization Dashboard

The VMware administrator uses the **vSAN Utilization** dashboard with the **vSAN Contention** dashboard to manage performance.

Design Consideration

The **vSAN Utilization** dashboard supports the **vSAN Contention** dashboard. Use it to identify vSAN clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

How to Use the Dashboard

- Clusters Utilization.

  - It lists all the vSAN clusters, sorted by the least performing.

- Select a vSAN cluster from the clusters utilization table.

  - All the health charts show the KPI of selected cluster.

- Disk Groups

  - It lists all the vSAN clusters, sorted by the least performing.

- Select a Disk Group from the disk groups table.

  - All the health charts show the KPI of selected cluster.

Points to Note

- The **vSAN Utilization** dashboard complements the **vSAN Contention**. For more information, see the points to note in the vSAN Contention Dashboard.

vSAN File Services

The VMware administrator uses the **vSAN File Services** dashboard to monitor the file services running in their vSAN environment.

Design Considerations

This dashboard is designed to complement the vSAN file services management provided by the vCenter Server. The vCenter Server is more of an administrative tool, while vRealize Operations Manager is more of an operations tool. Each tool performs their specific functions and does not duplicate information.

How to Use the Dashboard

- File Shares by Used Space and Latency.

  - Review the file shares by used space and latency heat map.

  - It shows all the file shares in your environment.

  - The greater the use (consumption), the greater the box, so you can easily see the most consumed ones.

  - The file shares are colored by latency. You must watch out for boxes with red color.

- vSAN Clusters with File Services enabled.

  - It lists all the vSAN clusters with file services enabled, giving a convenient view to see which clusters have these settings turned on.

- Select a vSAN cluster from the vSAN clusters with file services enabled table.

  - The file servers in the selected vSAN cluster are shown. When you select a file server, it filters the file shares list to show the file shares in the selected file server.

  - The file shares in the selected vSAN cluster are shown. Selecting a file share displays all the relevant KPI on the file share.

Points to Note

vSAN File Servers and vSAN File Shares are two new objects in vRealize Operations Management Pack for vSAN.

Dashboard Library

Deprecated Dashboards

Deprecated dashboards are kept intact and are not updated as the changes in the new predefined dashboards are substantial. Deprecated dashboards will be kept for at least one release. See the Release Notes for information about why the dashboards are deprecated.

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 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.

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 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.

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.

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.

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 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.

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.

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.

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.

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.

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.

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 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.

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.

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.

You can identify the vSAN encryption status at cluster levels.

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.

Executive Summary Dashboards

The requirements of the CIO, Head of Global Infrastructure, and IT Senior Management vary from the requirements of the technical teams. The **Executive Summary** dashboards provide an overall information on capacity and inventory in business terms.

These dashboards allow you to display problems related to budget and resource, and provide visibility to the senior management into the live environment. By doing this, you can prove the need for additional hardware. If there is wastage that has to be reclaimed, you can display where and how large the wastage is using these dashboards. vRealize Operations Manager provides two example dashboards to get you started. As each executive might have a unique requirement or preference, the dashboards can be customized accordingly. The five principles displayed in the following figure are used to design the **Executive Summary**
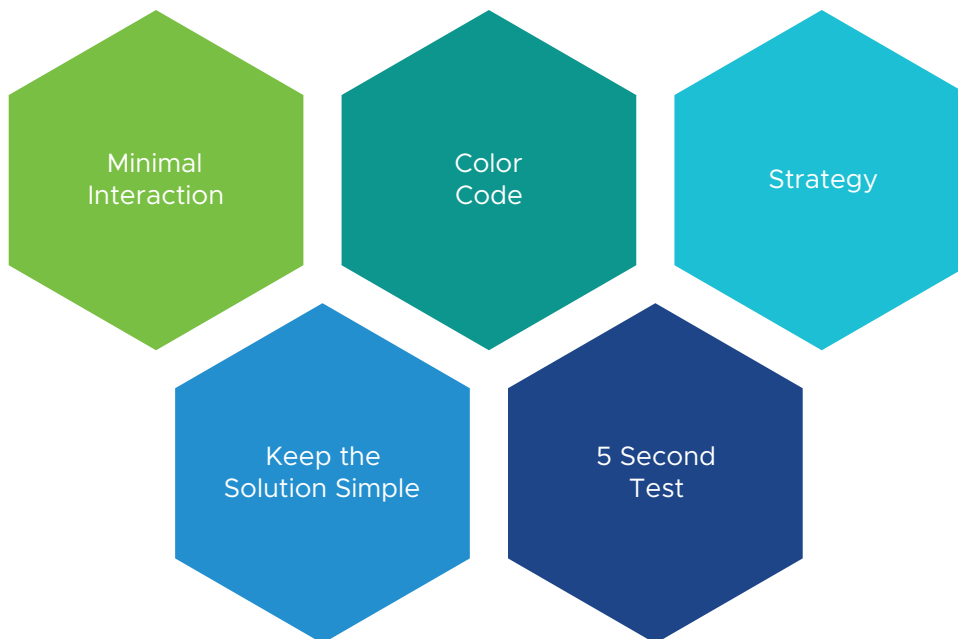
dashboards.

- Keeping the interaction, such as clicking, zooming, and sorting to a minimal.

- Use of color codes to have a user interface that is easy to understand.

- Each dashboard answers a specific question and the information is presented in business terms.

- Keep the solution simple and have a portal that is easy to access.

- Ensure that the dashboards are understood within five seconds.

Capacity Summary Dashboard

The **Capacity Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Inventory Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

How to Use the Dashboard

The **Capacity Summary** dashboard has two sections:

- The top section of the dashboard provides a summary at the vSphere World level.

  - The **VM Growth** widget displays the weekly average of the VM growth and provides holistic visibility of overall growth across all data centers for both running and powered off workloads. If an increase in the VM count is not accompanied by a corresponding increase in utilization, these newly provisioned VMs are likely not yet used.

  - The **Overcommit Ratio** widget highlights the efficiency gained by vSphere virtualization running multiple workloads on a shared infrastructure. Overcommitment has to be further reviewed along with elevated resource contention to understand the impact of performance on VMs competing for resources. In general, Overcommit is required to be financially more economical than the public cloud. As a reference, AWS typically overcommits CPU 2:1 by counting the hyper-threading and does not overcommit memory.

    Note vRealize Operations Manager uses physical CPU Cores not Logical Cores (Hyper-threading) for all CPU-based capacity calculations.

- The bottom section of the dashboard enables drill down into individual compute or storage capacity.

  - Capacity is split into Compute (vSphere Clusters) and Storage (Datastores) views. The heat map displays capacity by size and color by time remaining. By selecting either Clusters or Datastores, you can further drill down understand the remaining capacity and time (in days).

Points to Note

- Capacity remaining is not displayed at the vSphere World level as it can be misleading, especially in a global, or large infrastructure. Clusters also tend to serve a different purpose and they are not interchangeable.

- If you are using both on-prem and external cloud, for example, VMware on AWS, consider splitting the dashboard into two columns.

Inventory Summary Dashboard

The **Inventory Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Capacity Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

How to Use the Dashboard

- The **Summary** widget provides a quick view of the key inventory number.

  - The scoreboard is interactive. This widget drives the eight pie charts that are placed at the bottom of the dashboard. Since all the information is at the vSphere World level, clicking any of them will display details of the total inventory.

- Select any data center from the **Datacenters** widget.

  - This widget drives Clusters and Datastores so that you can quickly view what you have in a given data center and related capacity.

  - For a small environment, the vSphere World is displayed so you can view all the VMs in the environment.

  - To sort by any of the columns in the table, click the column title.

- The eight charts in the dashboard provide details of the inventory. They are driven by **Datacenters**, **Compute**, **Storage**, and **Summary** widgets.

Points to Note

- Understand the relationship hierarchy in vSphere. For example, Compute (cluster) is not a parent of Storage (datastore), so logically it is not possible to display datastores in a cluster. Data center consists of compute (cluster), network (distributed switch), and storage (datastore).

- Datastores do not drive the pie chart. This is a known limitation in the View widget.

- If your senior management wants to view the largest VM in a given environment, add a Top-N widget to list the top 10 largest consumers so that CPU, memory, and the disk details are highlighted.

Network Operation Center

A dashboard projected on the large screen serves a different business purpose than a dashboard on your laptop or desktop. It is placed strategically because it displays a time sensitive

information. Dashboards complement alerts and cannot replace it. The five principles displayed in the following figure are used to design the predefined **Network Operation Center** dashboards.

- Keeping the interaction, such as clicking, zooming, and sorting to a minimal. Avoid having buttons, use of mouse or keyboard to view data.

- Use of color codes to have a user interface that is easy to understand.

- Displaying content that drives action. Display of live information as the focus is on immediate remediation. Problems that need immediate actions are displayed, for example, stop provisioning of new VM or take action on VMs that abuse the shared infrastructure.

- Display of problems that do not require immediate attention are avoided, for example, increase supply of infrastructure, such as adding hardware.

- Keep the display simple and have a portal that is easy to access.

- Dashboards are designed to display minimal and critical information only.

- Displays of numbers in percentage, with 0% being poor and 100% being perfect. To display utilisation, you can use the following markers:

  - 50% indicates good and balanced utilization. However, the ideal value is 75%

  - 0% indicates wastage

- 100% indicates high utilization

- Ensure that the dashboards are understood within five seconds.

Live! Cluster Performance Dashboard

The **Live! Cluster Performance** dashboard provides live information on whether the requests of the VMs are met by their underlying compute clusters. This dashboard focuses on CPU, Memory, and the performance of the clusters. Use this dashboard to view if there is any problem in meeting the demands of the VMs and if there is any unbalance within a cluster. The **Live! Cluster Performance** dashboard is the primary dashboard and it complements the **Live! Cluster Performance** dashboard which is the secondary dashboard. This secondary dashboard displays if the performance problem is caused by high utilization. The primary dashboard answers the question 'Is our IaaS performing?', while the secondary dashboard answers the question 'Is our IaaS working hard?'.

Design Consideration

The **Live! Cluster Performance** dashboard displays three heat maps. The heat maps complement each other and must be used together. The location of each cluster and ESXi hosts within those clusters is identical in all heat maps. The fixed positioning allows you to compare if the problem is caused by memory contention, CPU ready, or CPU co-stop.

The sizes of each cluster and ESXi hosts are constant. Variable sizing creates a distraction and can result in small boxes, making it difficult to read.

The focus of the performance is on the population and not on a single VM. This is not a single VM troubleshooting dashboard but a dashboard focusing on infra problem. As the infra counter is mathematically an aggregation of VM counters, you must have a right roll-up strategy. As the goal is to provide an early warning, do not use the average as a roll-up technique. Use the percentage of the population exceeding a threshold. The threshold is set to be stringent to receive an early warning.

How to Use the Dashboard

Review the heat maps, **Memory Contention**, **CPU Ready**, and **CPU Co-Stop** and see if there is any color other than green.

- Green indicates that almost 100% of the VMs have received the CPU and memory that was requested. The threshold is set such that if the 10% of the VM population does not receive the requested resources, then the heat map turns red.

- Red indicates an early warning. Stringent thresholds are used to enable proactive attention and remediation operations. The heat map can turn red because of the high standard that is applied even when there is no complaint from the VM owner yet.

- The light gray indicates that there is no VM running on the host and the metric is not computing.

View if there is any unbalance.

- There are two types of unbalance, cluster unbalance, and resource type unbalance.

- The ESXi hosts are grouped by the cluster, so that the unbalance within a cluster can be easily viewed. Cluster unbalance is a real possibility and it is best monitored and not just assumed.

- If the three heat maps are different, then there is a resource unbalance. For example, if the memory contention is mostly red, but the two CPU heat maps are green, it means you have an unbalance between memory and CPU.

- If a single ESXi host displays different color across the three heat maps, it indicates that there is an unbalance between the CPU and memory resources in the host.

For NOC Operator, drill-down by selecting one of the VMs on the heat map.

- The **Trends of Selected ESXi Host** widget will automatically display the performance counters. To hide any metric, click the name in the legend.

As part of the deployment, configure auto-rotate among the NOC dashboards. If you want to view one dashboard, then you can remove the vRealize Operations Manager menu by using the URL sharing feature. This makes the overall user interface presentable and allows you to focus on the dashboard.

Points to Note

- You can add Disk Latency if you have the screen real estate. Use the counter 'Percentage of Consumers facing Disk Latency (%)'. It is a part of a datastore object, not a cluster, as a VM in a cluster can have disks across multiple datastores. Organize this storage performance by data center and not by the cluster.

Live! Cluster Utilization Dashboard

The **Live! Cluster Utilization** dashboard complements the **Cluster Performance** dashboard. Use this dashboard to view the clusters that are working excessively and are close to their physical limit. This dashboard displays ESXi hosts that have CPU or memory saturation that can lead to performance issues for the VMs running on the host.

Design Considerations

This dashboard is designed to complement the **Live! Cluster Performance** dashboard and it shares design considerations.

How to Use the Dashboard

As this dashboard has an identical design with the **Live! Cluster Performance** dashboard, it has the same usage procedure. Unlike the heat maps in the **Live! Cluster Performance** dashboard, the three heat maps in this dashboard have a different scale, reflecting the different nature of the counters.

Logically, memory is a form of storage. It acts as a cache to disk as it is much faster. A high utilization is better, as it indicates that more data is being cached. The ideal situation is when ESXi host Consumed metric is red but ESXi host Ballooned metric is green. When Ballooned is red and Consumed is gray, it means that there was high pressure in the past but it is not there anymore. The reason the ballooned stays red is because the ballooned pages were never requested back.

The ballooned memory counter was selected over the swapped or compressed memory counters as it is a better leading indicator. Since all three can co-exist at the same time, they are displayed in the line chart. Ballooned is displayed in absolute amount and not as a percentage, because the higher the size, higher are the chances for it to impact a VM. If you feel using percentage is easier for your operations, create a super metric to translate the value.

The heat map displays Wastage by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

Analyze if the ESXi host is contributing. A light gray box indicates that the host is a part of the cluster but there is no utilization. It is possible for the host to be in the maintenance mode or is powered off.

Points to Note

- ESXi host chooses to swap over compression if the compression ratio is less than 4x.

- If the ESXi host's physical NIC is saturated in your environment, then you can add a Network Throughput heat map.

Live! Heavy Hitters Dashboard

The **Live! Heavy Hitters** dashboard helps you analyze the misuse of the shared infrastructure. This dashboard displays details of VMs misusing shared infrastructure and if that has caused performance problems to the other VMs. The shared infrastructure includes risks. The cause for excessive load might be attacks, for example, the denial of service, process runaway, or a mass activation of agents. The most demanding VM is the largest. If a handful of VMs is dominating the shared infrastructure, their collective size is displayed on the dashboard.

Design Considerations

See the Performance Dashboards page for common design considerations among all the dashboards for performance management.

In a shared environment, it is possible to have a victim-villain problem. In the heat map, the villain VM is the one with the largest box size, while the victim VM is the one with the red box. If a handful of VMs is dominating the shared infrastructure, their collective size will be highly visible on the dashboard.

How to Use the Dashboard

- The heat maps, Disk IOPS, Disk Throughput, Network Throughput, and CPU Demand displays the four different loads that can be excessive. The heat maps display the relative value and not the absolute value. A VM does not generate a high load in the absolute term just because it has a large configuration.

- Each heat map has its color threshold, reflecting the nature of the contention metrics used in each of them.

- For NOC Operator, drill-down by selecting one of the VMs on the heat map. All the four line charts are automatically displayed, enabling you to get a complete picture of the selected VM.

Points to Note

- Memory is not displayed as it is a form of storage. The memory counters are space utilization and not speed. Think of disk space instead of IOPS. It can cause a capacity problem on the shared ESXi host, but not performance problems to other VMs.

- In a large environment, it might be difficult to view a small victim VM. Consider having multiple dashboards and use them interchangeably.

Software Defined Wide Area Network Dashboard

The Software-Defined Wide Area Network (SD-WAN) dashboard allows you to configure and monitor the services related to VeloCloud and SD-WAN using vRealize Operations Manager. Using the SD-WAN dashboard, you can also collect the metrics for VeloCloud Orchestrator and VeloCloud Gateway.

By default the SD-WAN dashboards are disbaled, if you want to know how to enable them, see Manage Dashboards. The following services are discovered using the VeloCloud Orchestrator:

- Java Application

- VeloCloud Orchestrator

- Nginx

- ClickHouse

- MySQL

- Redis

- Network Time Protocol

The following services are discovered using VeloCloud Gateway:

- Network Time Protocol

- VeloCloud Gateway

Troubleshoot SD-WAN Dashboard

You can use the widgets in Troubleshoot SD-WAN dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN.

You can use the dashboard widget in several ways:

- **Troubleshoot Virtual Machine (VM)**: Use this widget to navigate to a specific VM and troubleshoot the issues.

- **Troubleshoot Orchestrator**: Use this widget to navigate to a specific orchestrator and troubleshoot the issues.

- **Troubleshoot Gateway**: Use this widget to navigate to a specific gateway and troubleshoot the issues.

- **Troubleshoot Application**: Use this widget to navigate to a specific application and troubleshoot the issues.

- **Relationship**: Use this widget to view the services and operating system associated with the VeloCloud Orchestrator.

- **Top Alerts**: Use this widget to view the top alerts associated with the SD-WAN.

Troubleshoot SD-WAN Gateway Dashboard

You can use the widgets in Troubleshoot SD-WAN Gateway dashboard to monitor and troubleshoot all the services and applications associated with the SD-WAN gateway.

You can use the dashboard widget in several ways:

- **Active Alerts on the Gateway**: Use this widget to view the active alerts for the gateway.

- **Health of Gateway Applications**: Use this widget to view the health status of the applications in the gateway.

- **Examine Operating System**: Use this widget to examine the operating system status.

- **Gateway Summary Status**: Use this widget to view the summary information for the gateway.

- **Gateway Process Status**: Use this widget to view the process information for the gateway.

- **Gateway Resource Metrics**: Use this widget to view the resource metrics associated with the gateway.

- **Parent Host**: Use this widget to view the parent host information.

- **Parent Cluster**: Use this widget to view the parent cluster information.

Troubleshoot SD-WAN Orchestrator Dashboard

You can use the widgets in Troubleshoot SD-WAN Orchestrator dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN Orchestrator.

You can use the dashboard widget in several ways:

- **Active Alerts on the Orchestrator**: Use this widget to view the active alerts for the Orchestrator.

- **Health of Orchestrator Applications**: Use this widget to view the health status of the applications in the gateway.

- **Examine Operating System**: Use this widget to examine the operating system status.

- **Examing MySQL**: Use this widget to examine the MySQL application.

- **Orchestrator Service Status**: Use this widget to view the service status of the Orchestrator.

- **Redis Status**: Use this widget to view the status of the Redis application.

- **API Check Status**: Use this widget to check the API status.

- **Nginx Status**: Use this widget to check the Nginx status.

- **Parent Host**: Use this widget to view the parent host information.

- **Parent Cluster**: Use this widget to view the parent cluster information.

vRealize Automation 7.x Dashboards

With the vRealize Automation 7.x dashboards, you can monitor and troubleshoot objects in your cloud infrastructure.

The following vRealize Automation 7.x 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.

- **Virtual Machine**: Use this widget to view VMs that belong to the deployment.

- **Configured Users**: Use this widget to view information about the user that the virtual machine belongs to.

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 Capacity Remaining**: Use this widget to view the capacity constrained for a tenant object.

- **Business Group Capacity Remaining**: Use this widget to view the capacity constrained for a business group object.

- **Reservation Capacity Remaining**: Use this widget to view the capacity constrained for a 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.

vRealize Automation 8.x Dashboards

With the vRealize Automation 8.x dashboards, you can monitor and troubleshoot objects in your cloud infrastructure.

The following vRealize Automation 8.x dashboards are added to the predefined vRealize Operations Manager dashboards:

- Cloud Automation Environment Overview

- Cloud Automation Project Cost Overview

- Cloud Automation Resource Consumption Overview

- Cloud Automation Top-N Dashboard

Cloud Automation Environment Overview

You can use the widgets in the Cloud Automation Environment Overview dashboard to view the environment details for the vCenter Cloud Zone objects. You can use the Cloud Automation Environment Overview dashboard to view the projects, deployments associated with the vCenter Cloud accounts.

You can use the dashboard widgets in several ways.

- **vCenter Cloud Zone List**: Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the cloud zone objects present in your environment.

- **Project List**: Use this widget to view the total blueprints, cloud zones, deployments, virtual machines, health, risk, efficiency details in your environment.

- **Top Alerts**: Use this widget to view the top alerts in your environment.

- **VM List**: Use this widget to view all the VM details in your environment.

- **Blueprint List**: Use this widget to view the blueprint objects in your environment.

- **Deployment List**: Use this widget to view the blueprint objects deployed in your environment.

Cloud Automation Project Cost Overview

You can use the widgets in the Cloud Automation Project Cost Overview dashboard to view the project cost associated with cloud zone objects present in your environment.

You can use the dashboard widgets in several ways.

- **Project Cost**: Use this widget to view the project wise cost for compute, storage, and additional resources associated with your cloud environment.

- **Total Cost Over Time**: Use this widget to view the cost of individual projects on a day to day basis.

- **Deployment Cost by Selected Project**: Use this widget to view the deployment cost for the selected project in your cloud environment.

Cloud Automation Resource Consumption Overview

You can use the widgets in the Cloud Automation Resource Consumption Overview dashboard to view the resources consumed by vRealize Automation 8.x on Cloud Accounts.

You can use the Cloud Automation Resource Consumption Overview dashboard widgets in several ways.

- **Cloud Account**: Use this widget to view all the attributes related to the cloud account.

- **Cloud Zone**: Use this widget to view all the attributes related to the cloud zones.

- **Project**: Use this widget to view all the project details associated with your cloud account.

- **Cluster List**: Use this widget to view all the details associated with the clusters in your account.

- **Cluster Utilization**: Use this widget to view the cluster utilization details for the cloud accounts.

- **Deployment Heat Map by Project**: Use this widget to view the heat map for each deployed project in your cloud environment.

- **Cloud Zone Capacity**: Use this widget to view the memory and storage capacity that is allocated, reserved, and free for each cloud zone object.

- **Cloud Zone Memory Trend**: Use this widget to view and analyze a seven-day trend for the memory allocated, reserved, and free for the cloud zone.

- **Cloud Zone Storage Trend**: Use this widget to view and analyze a seven-day trend for the storage allocated, reserved, and free for the cloud zone.

Cloud Automation Top-N Dashboard

You can use the widgets in the Cloud Automation Top-N dashboard to view the projects with most critical alerts, to view the blueprint with most deployments, and to view the deployments with the highest cost.

You can use the dashboard widgets in several ways.

- **Project with Most Critical Alerts**: Use this widget to view the projects which has most critical alerts.

- **Top Alerts**: Use this widget to view the top alerts for the projects in your cloud account.

- **Blueprints with Most Deployment**: Use this widget to view the blueprint which has maximum deployments for the cloud account.

- **Relationship**: Use this widget to analyze the relationship between blueprints and deployments, and deployment and cost.

- **Deployment with Highest Cost**: Use this widget to identify the most expensive deployment associated with your cloud account.

Service Discovery Dashboards

Using the service discovery dashboards, you can determine the inter-dependencies of virtual machines and the dependencies of each service in the respective virtual machines.

The following service discovery dashboards are added to the predefined vRealize Operations Manager dashboards:

- Service Distribution

- Service Relationships

- Service Visibility

- Virtual Machine Relationships

Service Distribution Dashboard

You can use the dashboard to view the distribution of different services in the selected data center, cluster, or a host system. You can also view known and unknown services including the category and distribution percentage across a vSphere resource.

You can use the dashboard widgets in several ways:

- **Inventory Item**: Use this widget to view a hierarchical representation of objects in the form of badges.

- **Known Services Distribution**: Use this widget to view different services discovered from a selected object.

- **Service Categories**: Use this widget to view the service categories that are discovered by selecting an object from the resource widget.

- **User Defined Services Distribution**: Use this widget to view a list of user-defined services.

Service Relationships Dashboard

You can use the dashboard to view properties of the service such as the install path, the ports used, and the version. You can also view the relationship between the services that run on other VMs.

You can use the dashboard widgets in several ways:

- **List of Services Discovered**: Use this widget to view the services that have been discovered.

- **Connections from the Selected Services**: Use this widget to view the relationship between the services and the other services running on the VMs.

- **Properties of the Selected Service**: Use this widget to view the properties of the selected services.

Service Visibility Dashboard

You can use the dashboard to view a list of VMs without service visibility and VMs with user-defined services after you select a vSphere object.

You can use the dashboard widgets in several ways:

- **Inventory Tree**: Use this widget to view a hierarchical representation of objects in the form of badges.

- **Virtual Machines without Service Visibility**: Use this widget to view information about services where discovery has failed.

- **Virtual Machines with User-Defined Services**: Use this widget to view a list of VMs where the user has defined such services.

Virtual Machine Relationships Dashboard

You can use the dashboard to view a list of VMs with service discovery details such as, status, method, incoming/outgoing connections, and protection groups. When you select a VM, the dashboard displays a list of discovered services on the VM, the relationships of the VMs with other VMs based on the relationships of the discovered service.

You can use the dashboard widgets in several ways:

- **List of virtual machines**: Use this widget to view all the VMs discovered by the vCenter Server.

- **Node relationship of the selected VM**: Use this widget to view the relationship between the objects.

- **List of Services running in the selected VM**: Use this widget to view all the properties of the selected VM.

- **Connections of Virtual Machines**: Use this widget to view the relationship between one or more VMs.

Inventory Dashboards

The three vSphere Inventory dashboards and workload management inventory dashboards cater to the compute, network, and storage aspects of your SDDC. Using these dashboards, you can navigate through the environment and view your inventory and their key metrics at a glance. The Network and Storage dashboards can be shared with the network and storage teams respectively, giving them the necessary visibility, and increasing the collaboration between teams.

vSphere inventory dashboards

The vSphere inventory dashboards are built specifically for each role, but they share a common design. They have a similar layout and are used in the same manner. This makes learning easier, especially in smaller environments where the same team manages the full environment.

These dashboards help you answer several key questions:

- What is the topology of your vSphere compute inventory?

- What is the topology of your vSphere storage inventory?

- What is the topology of your vSphere network inventory?

Workload Management Inventory Dashboard

This is a unified dashboard for the new workload management objects. It shows the relationships and KPIs for the workload management objects. For example, you can see the topology view from the Tanzu Kubernetes clusters to the physical infrastructure.

vSphere Compute Inventory Dashboard

You can use the vSphere Compute Inventory Dashboard to browse through the topology of your vSphere compute inventory which includes information related to vSphere world, vCenter Server, data center, clusters, hosts, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the clusters, ESXi hosts, and virtual machines associated with the object.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to an object in the environment.

- **Metrics**: View the metrics related to the object.

- **Clusters**: View the cluster functionality.

- **ESXi Hosts**: View the data related to the hosts.

- **Virtual Machines**: View VMs that belong to the object.

vSphere Network Inventory Dashboard

The vSphere Network Inventory Dashboard allows you to browse through the topology of your vSphere network inventory which includes information related to vSphere world, vCenter Server, data center, distributed vSwitches, distributed port groups, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the distributed vSwitches, distributed port groups, virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to the object in the environment.

- **Metrics**: View the metrics of the object.

- **Distributed vSwitches**: View details related to the distributed vSwitches.

- **Distributed Port Groups**: View data relevant to distributed port groups.

- **Virtual Machines**: View VMs that belong to the object.

vSphere Storage Inventory Dashboard

The vSphere Storage Inventory dashboard allows you to browse through the topology of your vSphere storage inventory which includes information related to vSphere world, vCenter Server, data center, datastore clusters, datastores, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the datastore clusters, datastores, and virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to the object in the environment.

- **Metrics**: View the metrics of the object.

- **Datastore Clusters**: View the datastore cluster functionality.

- **Datastores**: View the datastore functionality.

- **Virtual Machines**: View VMs that belong to the object.

Workload Management Inventory Dashboard

The Workload Management Inventory dashboards curates the Kubernetes inventory across all the Workload Management enabled vSphere environments and displays it here. This includes an end to end topology map showcasing the health of all the objectes along with upstream and downstream dependencies. Upon clicking any object in the relationship tree, the related inventorty of Supervisor Clusters, Namespaces, Pods, Developer Managed VMs and Tanzu Kubernetes clusters can be viewed and exported from this dashboard.

You can select an object type to view the properties and key metrics related to it.

You can use the dashboard widgets in several ways.

- **Environment Summary**: Provides a summary of the supervisor cluster and the child objects.

- **Relationships**: An interactive canvas where you can view the relationship between the different objects in the workload management inventory.

- **Properties**: View the properties related to the object in the environment.

- **Metrics**: View the metrics of the object.

- **Supervisor Clusters**: View the supervisor cluster functionality.

- **Tanzu Kubernetes cluster**: View the Tanzu Kubernetes cluster functionality.

- **Virtual Machines**: View VMs that belong to the object.

- **vSphere Pods**: View information about vSphere Pods.

Microsoft Azure Dashboards

Use dashboards to monitor and troubleshoot Microsoft Azure issues in vRealize Operations Manager.

To access the dashboards, click **Dashboards** on the menu and click the dashboard names that start with Azure.

The following dashboards are available:

Dashboard Name	Purpose
Availability	View the availability of each Microsoft Azure service. Available services are green. Unavailable services are red and will be removed.
Inventory	View the adapter instance count in each resource group. Select a resource group to see a sparkline chart and the metrics for all the resources in the group. Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the inventory for the database. Note Metrics that are not collected or created are grayed out.
Optimization	View whether you are effectively using Microsoft Azure services. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
Virtual Machine	Select a virtual machine to view its scoreboard, property list, object relationship with resource group, and CPU usage and forecasting. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
SQL Database	Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the scoreboard, object relationship, and CPU usage for the database. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
Load Balancer	Select a load balancer to view its scoreboard, object relationship, and data path availability. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.

AWS Dashboards

Dashboards provide the user interface you use to monitor and troubleshoot Amazon Web Services problems in vRealize Operations Manager.

You can access the dashboards by selecting **Dashboards**, and then selecting **AWS**.

Table 9-1. AWS Dashboards

Dashboard Name	Purpose
AWS Alerts	The Alerts dashboard reports system-generated performance information for Amazon Web Services. In vRealize Operations Manager 5.8 and later, the dashboard also displays alerts received from Amazon Web Services Cloudwatch.
AWS ASG Utilization	Use the Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric. ASG metrics are not collected by default. You must enable them when creating the group. This applies only to the metrics belonging directly to the auto scale group, for example GroupDesiredCapacity. It does not apply to the aggregate instance metrics for the ASG, for example Instance Aggregate CPU Utilization.
AWS Disk Space	Use the Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default. For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html
AWS Instance Heatmap	Use the Instance Heatmap to monitor CPU/Disk/Network metric elements and identify instances that perform poorly.
AWS Instance Utilization	Use to identify which EC2 instances have high use across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by making adjustments to EC2 instances.

Table 9-1. AWS Dashboards (continued)

Dashboard Name	Purpose
AWS Troubleshooting	This dashboard is most helpful when someone calls in with a problem and you know which device they are using. You can search for that type of device or the specific device, if you know the name. When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas. You can use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another.
AWS Volume Performance	Use the Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations.
AWS Availability	Use this dashboard to view the availability of each AWS service.
AWS Inventory	Use this dashboard to view the count of each AWS service instance in each region.
AWS Optimization	Use this dashboard to view if you are effectively using AWS services.

Table 9-2. AWS - All Other Dashboards

Dashboard Name	Purpose
AWS Services ■ CloudFormation Stacks ■ Compute: EC2 ■ Compute: Elastic Containers ■ Compute: Lambda Functions ■ Database: Dynamo ■ Database: Elasticache ■ Database: RDS ■ Database: Redshift ■ Desktop: Workspaces ■ Network: Load Balancers ■ Network: VPS ■ Simple Queue Services ■ Storage	Select **AWS Services** and then select a dashboard to view a specific service-related information.

AWS Instance Utilization Dashboard

Use the AWS Instance Utilization dashboard to identify which EC2 instances have a high usage across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by adjusting the EC2 instances.

For example, you might determine that you need to resize the EC2 instance to make it larger or smaller.

You most often use this dashboard to troubleshoot issues with the listed metrics based on a support request from a user.

You can also identify which EC2 instances have been running for the longest and shortest amount of time. Then, you can use that information to determine whether EC2 instances can be decommissioned, or discover instances that have been added and need to be tracked in inventory.

Memory metrics require that you implement an add-on for each EC2 instance. These add-ons cost extra, and are not included by default.

AWS Auto Scaling Group Dashboard

Use the AWS Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric.

AWS Troubleshooting Dashboard

When a user calls in with a problem and you know the name of the device they are using, can search for that type of device or the specific device and use the AWS Troubleshooting dashboard to get an overview of the system functionality.

When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas.

Use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another.

There is a suggested flow to using the widgets in this dashboard.

1   Start with only the AWS Object widget open, and find the item you want to inspect.

2   Select the item, then expand the AWS Relationship widget to view the item's status.

3   Select one or all the related objects, then view the Ordered Symptoms, Interesting Metrics, and Mash-up.

4   Optionally, drag widgets into a new configuration if it makes it easier for you to compare information that is meaningful to you.

5   Examine the list of ordered symptoms and determine which of these events, in the given order might cause the problem to occur.

AWS Instance Heatmap Dashboard

Use the AWS Instance Heatmap dashboard to monitor CPU/Disk/Network metric elements and identify instances that perform poorly.

You can use the Troubleshooting dashboard to find more detail, and research the root cause of issues. Then you can view the specific object instance to identify faulty processes and take a corrective action.

AWS Volume Performance Dashboard

Use the AWS Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations. When you identify the EC2 instance that generates the load, use the Troubleshooting dashboard to investigate further.

AWS Disk Space Dashboard

Use the AWS Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default.

For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html.

AWS Alerts Dashboard

The AWS Alerts dashboard reports system-generated performance information for Amazon Web Services. In vRealize Operations Manager 6.6 and later, the dashboard also displays alerts received from Amazon Web Services Cloud watch.

Dashboards in VMware Cloud on AWS

The **VMware Cloud on AWS** dashboards allow you to track the capacity, cost, and inventory overviews of the SDDCs. You can also track the virtual machines monitoring and the utilization and performance of these SDDCs.

VMC Capacity Dashboard

Use the **VMC Capacity** dashboard to view the capacity overview of each VMware Cloud on AWS SDDC. You can view the capacity of Clusters, Hosts, VMs, Datastores, and Disk groups.

Table 9-3. Widgets in VMC Capacity Dashboard

Widget	Description
VMC SDDC by Capacity Remaining %	Displays the SDDCs as cards that show the remaining capacity percentage.
VMC SDDC by Time Remaining %	Displays the SDDCs as cards that show the remaining time percentage.
VMC SDDC by Virtual Machine Remaining (based on avg VM profile)	Displays the SDDCs as cards that show the remaining number of virtual machines.

When you select one of the SDDC cards, the details of that SDDC are automatically populated in the widgets after the VMC SDDC by Virtual Machine Remaining (based on avg VM profile) widget.

Note The key kpis are color-coded to help in identifying capacity bottlenecks.

VMC Cost Overview Dashboard

Use the **VMC Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

Table 9-4. Widgets in VMware Cloud on AWS Dashboard

Widget	Description
Organization Cost Overview	Displays a list of organizations with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD).
Outstanding Expense Trend	Displays the outstanding expense trend of the organization selected in the Organization Cost Overview widget.
Total Expense Trend (Monthly)	Displays the total monthly expenses trend of the organization selected in the Organization Cost Overview widget.
Commit Expense Trend (Monthly)	Displays the committed monthly expense trend of the organization selected in the Organization Cost Overview widget.
On-Demand Expense Trend (Monthly)	Displays the on-demand monthly expenses trend of the organization selected in the Organization Cost Overview widget.
Purchase History	Displays the bill line items/purchases from the available bills.
Currency Information	Represents the metrics currency unit set in this management pack account.

Note The YTD metric is an aggregation from the beginning of the calender year, until the last available bills.

VMC Inventory Dashboard

Use the **VMC Inventory** dashboard to view the inventory overview of all the SDDCs configured in VMware Cloud on AWS.

Widgets in VMC Inventory Dashboard

VMC SDDCs: displays the SDDCs as cards that show the number of virtual machines running in the SDDC. The SDDC card also shows a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that SDDC, the SDDC card indicates this by changing colors.

When you select one of the SDDC cards, the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs with key configuration details of that SDDC are populated in the widgets after the VMC SDDCs widget.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

VMC Management VM Monitoring Dashboard

Use the **VMC Management VM Monitoring** dashboard to monitor the utilization and performance of the key management VMs running in your SDDC. This dashboard ensures that the management components (such as vCenter and NSX) are not facing any resource bottlenecks from the CPU, memory, network, and storage perspectives.

Table 9-5. Widgets in VMC Management VM Monitoring Dashboard

Widget	Description
CPU Usage & Performance	Displays the list of all the management components in each SDDC with key CPU utilization and performance KPIs. Select a management VM to see the usage and performance trends of all the CPU cores.
Memory Usage & Performance	Displays the list of all the management components in each SDDC with key Memory utilization and performance KPIs. Select a management VM to see the memory usage and performance trends.
Network Usage & Performance	Displays the list of all the management components in each SDDC with key Network utilization and performance KPIs. Select a management VM to see the memory usage and performance trends.
Storage Usage & Performance	Displays the list of all the management components in each SDDC with key storage utilization and performance KPIs. Select a management VM to see the network usage and performance trends.

VMC Utilization and Performance Dashboard

Use the **VMC Utilization and Performance** dashboard to view the utilization and performance overview of each SDDC based on heavy hitter VMs and impacted VMs over the last 30 days. This dashboard helps you in finding the VMs in your environment that are negatively impacting the capacity or performance from a CPU, memory, storage, or network perspective.

Widgets in VMC Utilization and Performance Dashboard

List of VMC SDDCs: displays the list of all the SDDCs with aggregate CPU, memory, and storage utilization with 95th percentile and maximum values over the last 30 days.

When you select one of the SDDC from the List of VMC SDDCs widget, you can see the list of top VMs that are consuming compute, network & storage resources in that SDDC. The widgets after that show the compute (CPU & memory) utilization and performance analysis, network, storage, and utilization and performance analysis.

Each section in the dashboard is based on the last 30 days data with 95th percentile transformation which is configurable to Max, Average, Current, Standard Deviation, or other mathematical transformations.

Dashboards in NSX-T Management Pack

The **NSX-T Main** dashboard provides an overview of the network objects. It displays the topology of a selected object, how it connects to the elements in the network, and a view of related alerts.

Table 9-6. Widgets in NSX-T Main Dashboard

Widget	Description
NSX-T Instances	Displays the list of environments that are being monitored. When you select an environment in this widget, the other widgets in the **NSX-T Main** dashboard display data for the selected adapter.
Environment Overview	Displays a top-level view of the selected environment and following key components. ■ NSX-T Manager ■ Controller Node ■ Logical Router ■ Logical Switch ■ Load Balancer Virtual Server ■ Transport Zone
Top Alerts	Displays all the open alerts for the selected object in the **Environment Overview** widget.
Topology Graph	Displays the topology of the selected object in the **Environment Overview** widget.

NSX-T Configmax Metrics

The **NSX-T Configmax Metrics** dashboard provides an overview of all the configuration maximum metrics in all the NSX-T instances.

Table 9-7. Widgets in NSX-T Configmax Metrics Dashboard

Widget	Description
Select an adapter instance	Displays the list of all the NSX-T and NSX-T on VMC instances. When you select an instance in this widget, the other widgets in the **NSX-T Configmax Metrics** dashboard display data for the selected instance.
Relationship view	Displays the objects hierarchy for the instance selected in the **Select an adapter instance** widget. Only the objects with configuration maximum metrics are shown in the relationship view.
Select object from relationship view for the configmax metric	Displays all the configmax metrics for the selected object in the **Relationship View** widget.
Trend View	Displays all the MGW, CGW, and Distributed firewall section rule trends of the instance selected in the **Select an adapter instance** widget.
	Note The Trend View widget loads the trends only for the firewall sections object on VMware Cloud on AWS instances.

Configuring Administration Settings

<div style="text-align: right; font-size: 3em;">10</div>

After vRealize Operations ManagervRealize 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:

- vRealize Operations Manager License Keys

- vRealize Operations Manager License Groups

- vRealize Operations Manager Maintenance Schedules

- Manage Maintenance Schedules

- Managing Users and Access Control in vRealize Operations Manager

- vRealize Operations Manager Passwords and Certificates

- Modifying Global Settings

- Managing Content

- Transfer Ownership of Dashboards and Report Schedules

- vRealize Operations Manager Logs for Product UI

- Create a vRealize Operations Manager Support Bundle

- vRealize Operations Manager Dynamic Thresholds

- vRealize Operations Manager Adapter Redescribe

- Customizing Icons

- Allocate More Virtual Memory to vRealize Operations Manager

vRealize Operations Manager License Keys

To activate vRealize Operations Manager monitoring , you add licenses at installation or later. You track licenses so that you know what vRealize Operations Manager can monitor and when your licenses expire. A new license key is required for vRealize Operations Manager 7.0 and later versions. All license keys except vSOM Enterprise Plus and its add-ons are invalidated. The product works in evaluation mode until a new valid license key is installed. After you log in to

the user interface of vRealize Operations Manager, if you see that you are using an evaluation license, consider applying for a new license before the end of the 60-day evaluation period.

You can obtain the new license keys from the MyVMware portal.

Note If you added new licenses when you upgraded to vRealize Operations Manager 7.0, you can skip this step. However, if you have deployed a new instance of vRealize Operations Manager 8.x, you must install a new license.

How License Keys Work

License keys activate the solution or product and are available in varying levels. Higher levels typically allow vRealize Operations Manager to monitor more objects.

Where You Find the License Keys

1   In the menu, click **Administration**, and in the left pane click **Management > Licensing**.

2   Click the **License Keys** tab.

License Key Options

The options include toolbar and data grid options.

Click **Add** or click the **Hortizontal Ellipses** to refresh or remove license keys.

Table 10-1. License Key Toolbar Options

Option	Description
Add	Select a solution or product, and then enter and validate a license key for it.
Delete	Remove a license key.
Refresh License Usage	Update the list of keys.

Use the data grid options to view item details.

Table 10-2. License Key Data Grid Options

Option	Description
Product or Solution	Name of the product or solution associated with the key.
License Type	Level of the license. To view the license edition, click the [icon] icon, and then click **About**. The About vRealize Operations Manager dialog box opens. You can view the version no and the license edition that is in use.
License Capacity	Number of objects that the license allows the product to monitor.

Table 10-2. License Key Data Grid Options (continued)

Option	Description
License Usage	Number of monitored objects that count against the capacity. If you have an unlimited capacity, this number is zero (0).
Status	Indicates whether the license is valid.
Expiry	Date and time when the license expires.
License Information (below)	Details for the selected license key.
Overview	Solution or product, expiration, capacity, type, and use of the selected license key.
Associated License Groups	License groups that this key is a member of, and the number of objects in the groups.

vRealize Operations Manager License Groups

Like other vRealize Operations Manager groups, you create a license group of objects as a way of gathering those objects for data collection. In this case, you are associating the objects with a product license.

How License Groups Work

License groups require that you select one or more keys that you already added for solution or product activation, and add objects as members to a custom group for those licenses. You might, for example, want to add objects into groups that are associated with a particular level of license key, and monitor or manage by level of key in order to control licensing costs.

Where You Find the License Groups

1   In the menu, click **Administration**, and then in the left pane click **Management > Licensing**.

2   Click the **License Groups** tab.

License Groups

vCloud Suite

Host CPU-based licenses applied to an object type "Host system" for a given set of clusters. When you apply a CPU license to a group containing Hosts, the VMs on the Hosts will still show "License is invalid" watermark.

VM Licenses

VM based licenses applied to an object type "Virtual Machine" for all other VMs except those on hosts licensed with vCloud Suite. When you apply a VM license key to Virtual Machines, the Hosts on which those VMs run will still show the "License is invalid" watermark.

Note In vRealize Operations Manager, it is possible to mix Operating System Instance (OSI) and CPU based licenses. By mixing difference kind of licenses, you will need to perform extra configurations, like creating separate license groups for each type of license keys (one for CPU and one for OSI (VM)). It is recommended that you use non overlapping exclusive Licensing Groups to have the best advantage when you mix OSI (VM) and CPU licensing.

However, in vRealize Operations Manager you cannot mix core and standard license with any other advanced and enterprise licenses.

Dynamic

Use dynamic membership criteria, not static "Always include/exclude" lists to avoid manual maintenance of license groups.

Note When the license is applied to the respective Object type of each License key, the related objects (parent or children) are also going to have to be included in membership for the License Group. License in invalid" watermark appears in vRealize Operations Manager 6.6 and later. For more information, see the following KB article 51556.

License Group Options

The license group options include toolbar and data grid options.

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 10-3. License Group Toolbar Options

Option	Description
Add	Launch a wizard to select licenses and objects, to create a new license group. You can also associate the license group with a monitoring policy.
Edit	Launch a wizard to select licenses and objects, to change a license group. You can also associate the license group with a monitoring policy.
Delete	Remove a license group.

Use the data grid options to view item details.

Table 10-4. License Group Data Grid Options

Option	Description
License Group	Name of the license group
Total Members	Number of objects in the license group
Licensable Usage	Number of objects in the group that count against the license in order to monitor them. If you have a license for unlimited object monitoring, this number is zero (0).
License Group Information (below)	Details for the selected license group
Overview	Name, license serial number, and number of keys associated with the selected license group
Members	List of objects associated with the selected license group

vRealize Operations Manager Maintenance Schedules

Maintenance schedules identify objects that are in maintenance mode at specific times, which prevents vRealize Operations Manager from showing misleading data based on those objects being offline or in other unusual states because of maintenance.

Many objects in the enterprise might be intentionally taken offline. For example, a server might be deactivated to update software. If vRealize Operations Manager collects metrics when an object is offline, it might generate incorrect anomalies and alerts that affect the data for setting dynamic thresholds for the object attributes. When an object is identified as being in maintenance mode, vRealize Operations Manager does not collect metrics from the object or generate anomalies or alerts for it. In addition, vRealize Operations Manager cancels any active symptoms and alerts for the object.

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object in maintenance mode from midnight until 3 a.m. each Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can manually put an object in maintenance mode, or take it out of maintenance mode, even it if has an assigned maintenance schedule

Note   When you perform maintenance operations, it is good practice to stop the End Point Operations Management agent and to restart it after the maintenance is complete to avoid unnecessary system overhead.

How Maintenance Schedules Work

Maintenance schedules require that you select the days and time-of-day when updates or other object maintenance occurs. Note that creating a maintenance schedule does not activate the schedule. A maintenance schedule must be part of a policy before the schedule can take effect.

Where You Find the Maintenance Schedules

In the menu, click **Administration**, and then in the left pane click **Configuration > Maintenance Schedules** .

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 10-5. Maintenance Schedule Toolbar Options

Option	Description
Add	Open a window in which you can select the maintenance schedule settings for a new schedule.
Edit	Change the maintenance schedule settings for an existing schedule.
Delete	Remove the selected maintenance schedule.

Manage Maintenance Schedules

Add or edit a maintenance schedule to take an object offline. vRealize Operations Manager does not collect data from an object that is offline.

Where You Find Manage Maintenance Schedules

1   In the menu, click **Administration**, and then in the left pane click **Configuration > Maintenance Schedules**.

2   Click the plus sign to add a maintenance schedule or the pencil to edit the selected object.

Table 10-6. Manage Maintenance Schedule Add or Edit Options

Option	Description
Schedule Name	Name that describes the maintenance schedule
Time Zone	Time zone in which you are currently located
Days	Number of days the maintenance period covers
Recurrence	Specify a maintenance schedule to run over a selected period ■ Once ■ Daily ■ Weekly ■ Monthly
Expire after	The number of times the schedule is run
Expire on	The date upon which the schedule stops running

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.

- 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.
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.
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.

Access Control in vRealize Operations Manager

Each user must have a unique account with one or more roles assigned to enforce a role-based security when they use vRealize Operations Manager. You create a user account, and assign the account to be a member of one or more user groups to allow the user to inherit the roles and objects associated with the user group.

Where You Find the Access Control Options

You can manage user accounts and their associated user groups, roles, and passwords.

In the menu, click **Administration**, and then click **Access > Access Control** .

Table 10-7. Access Control Tabs

Option	Description
User Accounts	Add, edit, remove, or import vRealize Operations Manager user accounts from an LDAP database, and manage user roles, their membership in groups, and the objects assigned for association with the user. Import user accounts from an LDAP database that resides on another machine. vCenter Server users who are logged in to vRealize Operations Manager, either logged in directly or through the vSphere Client, appear in the list of user accounts.
User Groups	Add, edit, or remove, or import vRealize Operations Manager user groups, update the members in a group and the associated objects that they can access. Import user groups from an LDAP database or a single sign-on database that resides on another machine. vRealize Operations Manager continuously synchronizes the user membership of imported LDAP user groups when the autosync option is enabled in the LDAP configuration.
Roles	For users to perform actions in vRealize Operations Manager, they must be assigned specific roles. With role-based access, when you assign a role to a user, you are determining not only what actions the user can perform in the system, but also the objects upon which those actions can be performed while holding the role. For example, to import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management.
Password Policy	Manage local user passwords, set the criteria for account lockout, password strength, and the password change policy settings.

Access Control: User Accounts Tab

You can add, edit, or remove vRealize Operations Manager user accounts, and import user accounts from an external LDAP database. With access control, you manage roles, the objects a user can access while assigned a specific role, and the membership in user groups.

Where You Manage User Accounts

In the menu, click **Administration**, and then click **Access > Access Control**.

Table 10-8. Access Control User Accounts Summary Grid

Summary Grid Options	Description

User Accounts toolbar	To manage user accounts, use the toolbar icons.
	
	■ Click the **Add** icon to add a user account, and provide the details for the user account in the Add User Account dialog box.
	■ Click the **Vertical Ellipses** to perform any one of the following actions:
	  ■ **Edit**. Edit the selected user account, and modify the details for the user group in the Edit User Account dialog box.
	  ■ **Delete**. Delete a user account.
	■ Click the **Horizontal Ellipses** and click **Import Users** to import a user account from an authentication source.

First Name	User's first name, created when you create the user account.
Last Name	User's last name, created when you create the user account.

Table 10-8. Access Control User Accounts Summary Grid (continued)

Summary Grid Options	Description
User Name	User name, without spaces, that will log in to vRealize Operations Manager
Email	User's email address, created when you create the user account.
Description	Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access privileges.
Imported	Indicates whether the user account is imported or not.
Source Type	Indicates whether the user account is a local user, or an external user who is integrated through an external authentication source, such as from LDAP, SSO, AD, OpenLDAP, vCenter Server.
Enabled	Indicates whether the user account is enabled to use vRealize Operations Manager features. An administrator can edit a user account to manually enable it, or disable it to prevent user access to vRealize Operations Manager.
Locked	Indicates whether vRealize Operations Manager has locked the user account. For example, a user account can get locked based on the password lockout policy, or if the user enters an incorrect password three times in the span of five minutes.
Access All Objects	Indicates whether the user account is allowed to access all the objects that are imported into the vRealize Operations Manager instance.

After you add a user account, use the Details grid to view and edit which user accounts are assigned to user groups, and view the permissions assigned to the user account.

Table 10-9. Access Control User Accounts Details Grid

Details Grid Options	Description

User Groups	Assigned user groups appear when you click a user in the summary grid. You can then view and modify which user groups the user is associated with.
	
	■ User Name: Identifies the user account. To change the user groups associated with the user account, click the **Edit** icon.
	
	  The **Choose Groups Membership** dialog box opens.
	  ■ Click the **All** tab to view all the available groups.
	  ■ Click the **Selected** tab to view the groups that the user account is part of.
	  ■ Click the **Unselected** tab to view the groups that the user account is not a part of.
	  ■ Use the **Search** field to search for specific groups.
	■ Members: Displays the number of users that are assigned to the user group.

Permissions	Permissions appear when you click a user in the summary grid, and click the **Permissions** tab in the Details grid. You can then view the roles assigned to the user, and object hierarchy details. ■ Role: Indicates the name of the role or roles assigned to the user. ■ Role Description: Displays the description entered for the role. ■ Object Hierarchy: Displays the name of the object hierarchy assigned to the user while holding this role. ■ Objects: Displays the number of objects included in the hierarchy that the user can access. ■ Association: Indicates if the role and objects are assigned to the selected user, or assigned to a user group to which the user belongs.

Modify User Accounts and Assign Groups and Permissions

You can add user accounts so that users can access the features of vRealize Operations Manager and certain objects in the environment. Or, modify user accounts to change their attributes, disable or lock the accounts, or require them to change their password. After you add user accounts, you can assign them to one or more user groups, and assign roles and objects to the account to specify the actions the user can perform and upon what objects. Assign the administrators role only to specific users who must access objects and perform actions in the entire environment.

Where You Add or Edit User Accounts

1 To add a user account, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**

2 In the **User Accounts** tab, click **Add**

3 To edit a user account, click the vertical ellipsis and select **Edit**.

Table 10-10. Add or Edit Users Accounts- User Details Page

User Details Options	Description
User Name	User name, without spaces to access the vRealize Operations Manager
Password	User's password to access the vRealize Operations Manager instance.
Confirm Password	Confirmation of the user's password.
First Name	User's first name, created when you create the user account.
Last Name	User's last name, created when you create the user account.
Email Address	User's email address, created when you create the user account.
Description	Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access rights.
Disable this user	Disable the user account so that a user cannot access the vRealize Operations Manager instance.
Account is locked out	Indicates that vRealize Operations Manager has locked the user account.
Require password change at next login	Enable users to change their password the next time they log in to the vRealize Operations Manager instance.

4   After you enter the user details, click **Next**.

Table 10-11. Add or Edit User Accounts - Assign Groups and Permissions page

Assign Groups Roles, and Objects Options	Description
Groups	Select or deselect the groups associated with the user account. To select or deselect all accounts, click the **Group Name** check box. You cannot add user accounts to groups that you imported from an LDAP database.
Objects	Roles determine which actions a user can perform in the system. Select a role from the **Select Role** drop-down menu, and then select the **Assign this role to the user** check box. You can associate more than one role with the user account.
	Select which objects the user can access when assigned this role.
	■ Select Object Hierarchies: Displays groups of objects. Select an object in this list to select all the objects in the hierarchy.
	■ Select Object: To select specific objects within the object hierarchy, click the down arrow to expand the list of objects. For example, expand the Adapter Instance hierarchy, and select one or more adapters.
	■ **Allow access to all objects in the system**: Select this check box to permit the user account access to all objects in the system.
	Note The roles and object permissions are interlinked when you assign more than one role to a user. For example, if the user has both, ReadOnly and PowerUser roles, the permissions associated with the PowerUser role will apply, because the PowerUser role includes the permissions associated with the ReadOnly role along with other permissions.
	If the user has a custom role and the PowerUser role and the permissions of the custom role are not included in the permissions of the PowerUser role, the permissions of both the roles are merged and applied to the user.
	The same rule (object permissions from different roles are merged) applies to the object hierarchies as well.

Import User Accounts

You can import user accounts so that users can access the features of vRealize Operations Manager and the objects in the environment. After you import user accounts, you can assign them to user groups and roles. You can also specify the objects users can access while using the assigned roles.

Where You Import User Accounts

1   To import user accounts, click **Administration**, and then in the left pane click **Access > Access Control**.

2   Click the horizontal ellipsis next to **Add** and then, click **Import Users**.

Table 10-12. Import Users from a LDAP Source

User Details Options	Description
Import From	LDAP host machine, Active Directory, or Other sources configured to import user accounts. ■ **Add** icon. Add an LDAP import source, and provide the information for the LDAP import source in the Add Source for User and Group Import dialog box. ■ **Edit** icon. Edit the selected LDAP import source, and modify the details in the Edit Source for User and Group Import dialog box.
User Name	Click **Change Credentials** to display the user name of the LDAP source credential used to import user accounts to the vRealize Operations Manager instance.
Password	Password for the LDAP source credential to import user accounts to the vRealize Operations Manager instance.
Search String	Enter a search string, and click **Search** to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

Table 10-13. Import Users from a VMware Identity Manager Source

User Details Options	Description
Import From	VMware Identity Manager configured as the source to import user accounts. ■ **Add** icon. Add a VMware Identity Manager import source, and provide the information for the VMware Identity Manager import source in the Add Source for User and Group Import dialog box. ■ **Edit** icon. Edit the selected VMware Identity Manager import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Search Prefix	Enter a search string, and click **Search** to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

Table 10-14. Import Users from a Single Sign On Source

User Details Options	Description
Import From	SSO source configured as the source to import user accounts. ■ **Add** icon. Add an SSO import source, and provide the information for the SSO import source in the Add Source for User and Group Import dialog box. ■ **Edit** icon. Edit the selected SSO import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Result Limit	Determines the number of users displayed.

Table 10-14. Import Users from a Single Sign On Source (continued)

User Details Options	Description
Search Prefix	Enter a search prefix, and click **Search** to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

3   After you enter the import users details, click **Next**.

Table 10-15. Import Users Accounts- Assign Groups and Permissions Page

Assign Groups Roles, and Objects Options	Description
Groups	Select or deselect the groups associated with the user account. To select or deselect all accounts, click the **Group Name** check box. You cannot add user accounts to groups imported from LDAP.
Objects	Select or deselect roles in the **Select Role** drop-down menu. When you have selected a role, click the **Assign this role to the user** check box. You can assign more than one role to a user account. Select which objects the user can access when assigned this role. ■ Select Object Hierarchies: Displays groups of objects. Select an object in this list to select all the objects in the hierarchy, ■ Select Object: To select specific objects within the object hierarchy, click the down arrow to expand the list of objects. For example, expand the Adapter Instance hierarchy, and select one or more adapters. ■ **Allow access to all objects in the system**: Select this check box to permit the user account access to all objects in the system.

Access Control: User Groups Tab

You can manage the user groups associated with the users and objects in your environment. You can import user groups from an LDAP database that resides on another machine, or from a single sign-on server.

Where You Manage User Groups

1   To manage user groups, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.

2   Click the **User Groups** tab.

Table 10-16. Access Control User Groups Summary Grid

Option	Description
User Groups toolbar	To manage user groups, use the toolbar icons. ■ Click the **Add** icon to add a user group, and provide the details for the user group in the Add User Group dialog box. ■ Click the **Vertical Ellipses** to perform any one of the following actions: 　■ **Edit**. Edit the selected user group, and modify the details for the user group in the Edit User Group dialog box. 　■ **Clone**. Clone a user group, and enter a name and description for the cloned user group. 　■ **Delete**. Delete a user group. ■ Click the **Horizontal Ellipses** and click **Import** to import a user group, and provide the details to import the user group in the Import User Groups dialog box.
Group Name	Name of the user group.
Description	Description of the group, indicating its purpose.
Members	Number of members in the group.
Group Type	Type of group, either a local user group or a group imported from LDAP.
Distinguished Name	Names for LDAP objects, such as domains and users.
Access All Objects	Indicates if the user group account is allowed to access all the objects that are imported into the vRealize Operations Manager instance.

After you select a user group in the summary grid, view details about associated users in the Details pane.

Table 10-17. Access Control User Groups Details Grid

Option	Description

User Accounts	Associated user accounts appear when you click a user group in the summary grid. You can then view or modify user accounts that are part of the selected group.
	
	■ User Name: Name of each user who is a member of the selected group. To change the user accounts associated with the user group, click the **Add** icon.
	
	The **Add Users to Group** dialog box opens.
	  ■ Click the **All** tab to view all the available user accounts.
	  ■ Click the **Selected** tab to view the user accounts that are part of the group.
	  ■ Click the **Unselected** tab to view the user accounts that are not a part of the group.
	  ■ Use the **Search** field to search for specific user accounts.
	■ First Name: First name of each user account in the group.
	■ Last Name: Last name of each user account in the group.
	
	You can remove a user from the group by selecting the user in the Details pane and clicking **Delete**

Permissions	View the permissions of the role associated with the user group. To add or remove roles, view only the selected or deselected roles, or search for a specific role, click the **Edit** icon. ■ Role Name: Indicates the roles assigned to the selected user group. ■ Role Description: Description for the selected user group, defined when you created the group. ■ Object Hierarchy: The names of the object hierarchies assigned to the group while holding a specific role. ■ Objects: The number of objects the user group can access within the selected hierarchy.

Add User Groups and Assign Members and Permissions

You can view and modify the details for user groups, including users, roles, and objects.

Where You Add User Groups

1 To add a user group, in the menu, click **Administration** and then click **Access > Access Control**.

2 Select the **User Groups** tab and then click the **Add** icon.

Table 10-18. Add or Edit User Group - Name and Description Page

Option	Description
Group Name	Name of the user group, either created manually, imported from a single sign-on server, or imported from an LDAP database that resides on another machine.
Description	Description of the user group, indicating its purpose.

3 After you enter the name and description, click **Next**

Table 10-19. Add or Edit User Group - Assign Members and Permissions Page

Option	Description
Members	Select the members associated with the user group.
Objects	Roles determine which actions users of the group can perform in the system. Select a role from the **Select Role** drop-down menu, and then select the **Assign this role to the user** check box. You can associate more than one role with the user group.
	Select which objects the users of the group can access when assigned this role.
	■ Select Object Hierarchies: Displays groups of objects. Select an object in this list to select all the objects in the hierarchy.
	■ Select Object: To select specific objects within the object hierarchy, click the down arrow to expand the list of objects. For example, expand the Adapter Instance hierarchy, and select one or more adapters.
	■ **Allow access to all objects in the system**: Select this check box to permit users of the group access to all objects in the system.
	Note The roles and object permissions are interlinked when you assign more than one role to a user. For example, if the user has both, ReadOnly and PowerUser roles, the permissions associated with the PowerUser role will apply. The PowerUser role includes the permissions associated with the ReadOnly role along with other permissions.
	If the user has a custom role and the PowerUser role and the permissions of the custom role are not included in the permissions of the PowerUser role. The permissions of both the roles are merged and applied to the user.
	The same rule (object permissions from different roles are merged) applies to the object hierarchies as well.

Import User Groups

You import user groups from a single sign-on server, VMware Identity Manager, Active Directory, or an LDAP database on another machine so that you can use those groups in vRealize Operations Manager.

Where You Import User Groups

1   To import a user group, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.

2   Select the **User Groups** tab and click the **Import Group** icon.

    The options displayed in the Import User Groups page depend upon the authentication source you select.

Table 10-20. Import User Groups Page - LDAP, Active Directory, and Others Sources

Option	Description
Import From	Host machine configured as the source to import the user groups. These options are displayed when the host machine of an LDAP, Active Directory, or Other source is selected.
User Name	User name of the source credential to import user groups to the vRealize Operations Manager instance.
Password	Password for the source credential to import user groups to the vRealize Operations Manager instance.
Search String	Invoke the search for user groups.
Advanced	Displays the advanced import settings. ■ Group Search Criteria. Search criteria to find LDAP groups. If not included, vRealize Operations Manager uses the default search parameters: `(|(objectClass=group)(objectClass=groupOfNames))` ■ Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, vRealize Operations Manager uses member by default. ■ User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form `(|(key1=value1)(key2=value2))`. If not included, vRealize Operations Manager searches for each user separately. This operation might take extra time. ■ Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, vRealize Operations Manager treats the member entry as a distinguished name. ■ LDAP Context Attributes. Attributes that vRealize Operations Manager applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as `java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse`.
Group Name	Displays the user groups found. Click the check box for each user group to import.

Table 10-21. Import User Groups Page - Single Sign On Source

Option	Description
Import From	Host machine configured as the source to import the user groups.
Domain Name	User name of the source credential to import user groups to the vRealize Operations Manager instance.
Result Limit	Determines the number of groups displayed.
Search Prefix	Enter a search prefix to narrow your search.
Group Name	Displays a list of user groups. Select the **Group Name** check box to import all the displayed user groups, or select the check box next to each user group that you want to import.

Table 10-22. Import User Groups from a VMware Identity Manager Source

User Details Options	Description
Import From	VMware Identity Manager configured as the source to import user groups. ■ **Add** icon. Add an VMware Identity Manager import source, and provide the information for the VMware Identity Manager import source in the Add Source for User and Group Import dialog box. ■ **Edit** icon. Edit the selected VMware Identity Manager import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Search Prefix	Enter a search string, and click **Search** to start the search for user groups.
User Name Summary grid	Lists the users available for import. Select the check box for each user group to import, or select the **Group Name** check box to import all groups. User groups that are already imported to vRealize Operations Manager do not appear in the list.

3   After you enter the import user group details, click **Next**.

Table 10-23. Import User Groups - Roles and Objects Page

Option	Description
Select Role	Displays available roles in a drop-down menu.
Assign this role to the group	Roles determine which actions users of the group can perform in the system. Select a role from the **Select Role** drop-down menu, and then select the **Assign this role to the user** check box. You can associate more than one role with the user group.
Select Object Hierarchies	Select which objects the users of the group can access when assigned this role. ■ Select Object Hierarchies: Displays groups of objects. Select an object in this list to select all the objects in the hierarchy, ■ Select Object: To select specific objects within the object hierarchy, click the down arrow to expand the list of objects. For example, expand the Adapter Instance hierarchy, and select one or more adapters. ■ **Allow access to all objects in the system**: Select this check box to permit users of the group access to all objects in the system.

Access Control: Roles Tab

You can assign users-specific roles to perform actions and view features and objects in vRealize Operations Manager. With role-based access, users can only perform the actions that their permissions allow.

Where You Manage User Roles

1   To manage user roles, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.

2   Click **Roles** tab.

You can view and edit details about a role, by selecting a role in the summary grid, and clicking the **Edit** icon in the Roles toolbar.

Table 10-24. Access Control Roles Summary Grid

Option	Description

Roles toolbar	To manage roles, use the toolbar icons.
	
	■ Click the **Add** icon. to add a user role, and provide the name and description for the role in the Create Role dialog box.
	
	■ Click the **Vertical Ellipses** to perform any one of the following actions:
	
	  ■ **Edit**. Edit the selected user role, and modify the details for the role in the Edit Role dialog box.
	
	  ■ **Clone**. Clone the selected user role
	
	  ■ **Delete**. Delete a user role.

Role Name	Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions.
Role Description	Description of the role, indicating its purpose.

You can view details for the user accounts and user groups associated with a selected role in the Details panes.

Table 10-25. Access Control Roles Details Panes

Option	Description
User Accounts	The users assigned to the selected role. The information in this pane is based on the data entered when you created the user, or imported with the user. ■ First Name. Indicates the first name of each user who is assigned this role. ■ Last Name. Indicates the last name of each user who is assigned this role. ■ User name , without spaces, that will log in to vRealize Operations Manager ■ Email. Indicates the email address for each user who is assigned this role.
User Groups	The user groups assigned the selected role. ■ Group Name: Name of each group that is associated with the selected role. ■ Members: Number of members in each group.
Permissions	Displays the permissions assigned to the role according to five categories: Administration, Alerts, Dashboards, Environment, and Home. Expand the tree of each category to view all the assigned permissions. You can edit the permissions assigned to the role by clicking the **Edit** icon. ■ Click the **Expand All** button to expand the trees of all three categories, and select the check boxes to apply permissions for the selected role. ■ To assign all the available permissions to the selected role, select the **Administrative Access - all permissions** check box.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`

- `Set CPU Count for VM Power Off Allowed`

- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

Access Control: Password Policy Tab

To ensure security in vRealize Operations Manager, you must manage user passwords. Determine the criteria used for account lockout, password strength, and the password change policy. When a user session becomes inactive for 30 minutes, the session times out, and the user must log in to vRealize Operations Manager again.

Where You Manage the Password Policy

1 To manage user roles, in the menu, click **Administration**, and then click **Access > Access Control**.

2 Click **Password Policy** tab.

Account Lockout

Indicates whether the account lockout is in effect, and indicates the number of login attempts allowed before the account is locked. The account lockout policy is enabled by default.

Password Strength

Indicates whether the policy that requires users to strengthen their password is in effect, and the minimum number of characters required to make a strong password. The password strength policy is enabled by default.

Password Change

Indicates whether the policy that requires users to change their password is in effect, how often the password expires, and whether users will receive a warning. The account password change policy is enabled by default.

Modify the Password Policy

You can modify the password policy by clicking **Edit**.

Table 10-26. Access Control Edit Password Policy Settings

Option	Description
Account Lockout	Modify the settings to lock user accounts. ■ Activate Account Lockout Policy. Enable the policy to lock user accounts. For a super administrator user, the account lockout policy is enabled by default and cannot be disabled. The super administrator user account is locked for approximately one hour, and then unlocked. ■ Number of failed login attempts before lockout. Indicates the number of tries that a user can attempt to log in to vRealize Operations Manager before their account is locked. The default number of tries is seven, and the time frame allowed for login is 45 seconds.
Password Strength	Modify the settings required for users to create strong passwords. ■ Activate Password Strength Policy. When selected, enables the policy to require users to strengthen their password. ■ Minimum password length. Indicates the number of characters required for user passwords. The default length is eight characters. ■ Passwords must contain numbers. Users must include a combination of letters and numbers. ■ Passwords must not match user names. To ensure security, users are not allowed to use their user name as their password. ■ Passwords must contain at least one uppercase and one lowercase letter. When selected, users must include one or more uppercase characters. ■ Passwords must contain special characters. When selected, users must include one or more special characters. Special characters include: !@#$%^&*+=
Password Change	Modify the settings required for users to change their password. ■ Activate Password Change Policy. Enable the policy to require users to change their password at specific intervals. ■ Passwords expire every 90 days. Users receive notification five days before the password expires. ■ Warn users 5 days prior to expiration. Indicate when to have vRealize Operations Manager notify users that their password will expire. The default is five days before their password expires.

Authentication Sources

vRealize Operations Manager uses authentication sources that enable you to import and authenticate users and user group information that reside on another machine: the Lightweight Directory Access Protocol (LDAP) platform-independent protocol, Active Directory, VMware Identity Manager, Single Sign-On, and Others.

Where You Manage Authentication Sources

To manage authentication sources, in the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.

Table 10-27. Authentication Sources Toolbar and Data Grid

Option	Description

Authentication Sources toolbar	To manage authentication sources, use the toolbar icons.
	
	■ **Add** icon: Add an authentication source, and provide the information for the source in the Add Source for User and Group Import dialog box.
	
	■ Click the **Vertical Ellipses** to perform any one of the following actions:
	
	  ■ **Edit**. Edit the selected authentication source, and modify the details in the Edit Source dialog box.
	
	  ■ **Delete**. Delete an authentication source.
	
	  ■ **Synchronize User Groups**. Synchronize users within the groups imported through the selected Active Directory or LDAP authentication source.

Source Display Name	Name that you assign to the authentication source.
Source Type	Indicates the type of directory services access technology to access the source machine where the authentication database of user accounts resides. Options include: ■ Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts. ■ Active Directory or Other: Specifies any other LDAP-based directory services, such as Novel or Open DJ, used to import user accounts from an LDAP database on a Linux Mac machine. ■ SSO SAML: An open-standard data format that enables Web browser single sign-on. ■ VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources.
Host	Name or IP address of the host machine where the user database resides.
Port	Port used for the import.
Base DN	Base distinguished name for the user search. vRealize Operations Manager locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.
Auto Synchronization	When selected, enables vRealize Operations Manager to map imported LDAP users to user groups.
Last Synchronized	Date and time that the synchronization last occurred.

Authentication Sources: Add Authentication Source for User and Group Import

When you import user account information that resides on another machine, you must define the criteria used to import the user accounts from the source machine.

Where You Add or Edit Authentication Sources

1   To add authentication sources, in the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.

2   Click **Add**.

3   To edit authentication sources, click **Edit**.

Table 10-28. Authentication Sources Add Source for User and Group Import

Option	Description
Source Display Name	Name that you assign to the authentication source.
Source Type **Note** The option you select in the **Source Type** drop-down box, determines the options available in this dialog box.	Indicates the type of directory services access technology to access the source machine where the database of user accounts resides. There are two types of databases: LDAP and single sign-on. Options include: ■ SSO SAML: An XML-based standard for a web browser single sign-on that enables users to perform single sign-on to multiple applications. ■ Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts. ■ Other: Specifies any other LDAP-based directory services, such as Novel or OpenDJ, used to import user accounts from an LDAP database on a Linux Mac machine. ■ VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources.

Table 10-29. Authentication Sources Add Source for User and Group Import - Options Available When SSO SAML Is Selected.

Name	Description
Host	Name or IP address of the host machine where the single sign-on user server resides.
Port	The single sign-on listening port. By default this is set to 443.
User Name	Name of the user account that can log in to the single sign-on host machine.
Password	Password of the user account that can log in to the single sign-on host machine.

Table 10-29. Authentication Sources Add Source for User and Group Import - Options Available When SSO SAML Is Selected. (continued)

Name	Description
Grant administrator role to vRealize Operations Manager for future configuration?	When you create a single sign-on source, a new vRealize Operations Manager user account is created on the single sign-on server. ■ Select **Yes**, to grant vRealize Operations Manager an administrative role so that it can be used to configure the SSO source if changes are made to the vRealize Operations Manager setup. ■ If you select **No** and the vRealize Operations Manager setup is changed, SSO users will not be able to log in until you re-register the SSO source.
Automatically redirect to vRealize Operations single sign-on URL?	After you have configured a single sign-on source, users are redirected to the vCenter SSO server. ■ Select **Yes**, to redirect users to the single sign-on server for authentication. ■ If you select **No** users must sign in through the vRealize Operations Manager login page.
Import single sign-on user groups after adding the current source?	When you have set up a single sign-on source, you import users and user groups into vRealize Operations Manager so that single sign-on users can access the system with their single sign-on permissions. ■ If you select **Yes**, the wizard directs you to the Import User Groups page so that you can import user groups when you have finished setting up the SSO source. ■ If you want to import user accounts, or user groups at a later stage, select **No**.
Advanced	If your system uses a load balancer, enter the IP address of the load balancer.
Test	Tests whether the host machine can be reached with the credentials provided.

Table 10-30. Authentication Sources Add Source for User and Group Import - Options Available When **Open LDAP, Active Directory,** and **Other** Are Selected.

Option	Description
Integration Mode Basic settings	Applies basic settings to integrate the LDAP import source with the instance of vRealize Operations Manager.
	Use Basic integration mode to have vRealize Operations Manager discover the host machine where the LDAP database resides, and set the base distinguished name (Base DN) used to search for users. You provide the name of the domain and the subdomain, which vRealize Operations Manager uses to populate the Host and Base DN details, and the name and password of the user who can log in to the LDAP host machine.
	In Basic mode, vRealize Operations Manager attempts to fetch the host and port from the DNS server, and obtain the Global Catalog and domain controllers for the domain, with preference given to SSL/TLS-enabled servers.
	■ Domain/Subdomain. Domain information for the LDAP user account.
	■ Use SSL/TLS. When selected, vRealize Operations Manager uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, vRealize Operations Manager prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds.
	■ If Active Directory uses a self-signed certificate, then the certificate should contain the **Subject Alternative Name** field. vRealize Operations Manager can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the **Subject Alternative Name** field matches the address of the domain controller on which the certificate is used.
	■ User Name. Name of the user account that can log in to the LDAP host machine.
	■ Reset Password. Reset the password of the user account that can log in to the LDAP host machine.
	■ Automatically synchronize user membership for configured groups. When selected, enables vRealize Operations Manager to map imported LDAP users to user groups.
	■ Host. Name or IP address of the host machine where the LDAP user database resides.
	■ Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS.
	■ Base DN. Base distinguished name for the user search. vRealize Operations Manager locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.
	■ Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is *userPrincipalName*.
Integration Mode Advanced settings	Applies advanced settings to integrate the LDAP import source with the instance of vRealize Operations Manager.

Table 10-30. Authentication Sources Add Source for User and Group Import - Options Available When Open LDAP, Active Directory, and Other Are Selected. (continued)

Option	Description
	Use Advanced integration mode to manually provide the host name and base distinguished name (Base DN) to have vRealize Operations Manager import users. You provide the name and password of the user who can log in to the LDAP host machine.
	■ Host. Name or IP address of the host machine where the LDAP user database resides.
	■ Use SSL/TLS. When selected, vRealize Operations Manager uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, vRealize Operations Manager prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds.
	■ If Active Directory uses a self-signed certificate, then the certificate should contain the **Subject Alternative Name** field. vRealize Operations Manager can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the **Subject Alternative Name** field matches the address of the domain controller on which the certificate is used.
	■ Base DN. Base distinguished name for the user search. vRealize Operations Manager will locate only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.
	■ User Name. Name of the user account that can log in to the LDAP host machine.
	■ Reset Password. Reset the password of the user account that can log in to the LDAP host machine.
	■ Automatically synchronize user membership for configured groups. When selected, enables vRealize Operations Manager to map imported LDAP users to user groups.
	■ Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is *userPrincipalName*.
	■ Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS.

Table 10-30. Authentication Sources Add Source for User and Group Import - Options Available When Open LDAP, Active Directory, and Other Are Selected. (continued)

Option	Description
Search Criteria	Displays the search criteria settings. Although vRealize Operations Manager populates part of the search criteria, an Administrator must verify the settings to ensure that the settings are correct according to the properties of the LDAP type. ■ Group Search Criteria. Search criteria to find LDAP groups. If not included, vRealize Operations Manager uses the default search parameters: `(|(objectClass=group)(objectClass=groupOfNames))` ■ Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, vRealize Operations Manager uses member by default. ■ User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form `(|(key1=value1)(key2=value2))`. If not included, vRealize Operations Manager searches for each user separately. This operation might take extra time. ■ Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, vRealize Operations Manager treats the member entry as a distinguished name. ■ LDAP Context Attributes. Attributes that vRealize Operations Manager applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as `java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse`.
Test	Tests whether the host machine can be reached, with the credentials provided. Although a test of the connection is successful, users who use the search feature must have read permissions in the LDAP source. This test does not verify the accuracy of the Base DN or Common Name entries.

Table 10-31. Authentication Sources Add Source for User and Group Import - Options Available When VMware Identity Manager Is Selected.

Option	Description
Host	Name or IP address of the VMware Identity Manager machine where the single sign-on user server resides.
Port	The single sign-on listening port. By default this is set to 443.
Tenant	This is an optional field.
User name	VMware Identity Manager system-domain tenant administrator user name.
Password	Password of the VMware Identity Manager system-domain tenant administrator.

Table 10-31. Authentication Sources Add Source for User and Group Import - Options Available When VMware Identity Manager Is Selected. (continued)

Option	Description
Redirect IP/ FQDN	This is the IP address of vRealize Operations Manager node where a user is redirected after a successful authentication from VMware Identity Manager. By default, this is the IP address of the vRealize Operations Manager primary node.
	Note When the primary replica becomes the primary node on vRealize Operations Manager, then vRealize Operations Manager administrator has to manually edit the IP address and set it to the IP address of the current primary node.
Test	Tests whether the VMware Identity Manager machine can be reached, with the credentials provided.

Audit Users and the Environment in vRealize Operations Manager

At times, you might need to provide documentation as an 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 overlayed so that you can visualize relationships and the cause of the events. Events can include login attempts, system start up 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.

User Activity Audit

The user activity report helps you understand the scope of user activities in your vRealize Operations Manager instance, such as when users logged in, actions they took on clusters and nodes, changes they made to system passwords, when they activated certificates, and when they logged out.

Where You Audit User Activity

To audit user activity, in the menu, click **Administration**, and then in the left pane click **History > Audit**. The activities that users performed in the environment appear on the page.

Table 10-32. User Activity Audit Actions

Option	Description
Download	Download the user activity audit information to a report in PDF or XLS format.
Configure	Configure the settings to send the user activity log to an external syslog server to meet security auditing requirements. - Output log to external syslog server. When selected, vRealize Operations Manager sends the log to a separate server machine. - IP Address or Host Name. Identification for the syslog server. - Port. vRealize Operations Manager port used to send the audit information to the external server.

Table 10-32. User Activity Audit Actions (continued)

Option	Description
Date Range	Display the list of user activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times.
Starting Line	Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end.
Number of Lines	Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0.
Filter	Filters the data according to User ID, User Name, Auth Source, Session, Message, and Category.

User Permissions Audit

A user permissions audit report provides an overview of the local users and LDAP imported users in your vRealize Operations Manager instance, and a list of groups to which each user belongs. This report helps you understand the scope of the user accounts and their roles, access groups, and access privileges in your environment.

The report displays the access group associated with each local user and LDAP imported user and the access privileges granted to the user in each access group. This report does not include vCenter Server users, roles, or privileges.

The report displays the access group associated with each local user and the access privileges granted to the user in each access group. This report does not include vCenter Server users, roles, or privileges.

When a user is a member of a specific user group, the associated access group could provide the user with access to configuration, dashboards, and templates, or to specific navigation areas in the user interface such as Administration. The access rights associated with the access group include actions for each access group, such as the ability to add, edit, or delete dashboards, or to view, configure, or manage objects.

Where You Audit User Permissions

1   To audit user permissions, in the menu, click **Administration**, and then in the left pane click **History > Audit**.

2   Click the **User Permissions Audit** tab.

The permissions assigned to users, and their associated access groups and access privileges, appear on the page.

Table 10-33. User Permissions Audit Actions

Option	Description
Download	Download the user permissions audit information to a report in PDF or XLS format.

System Audit for vRealize Operations Manager

A system audit report provides an overview of the counts of objects, metrics, super metrics, applications, and custom groups in your vRealize Operations Manager instance. This report can help you understand the scale of your environment.

The system audit report displays the types and number of objects that vRealize Operations Manager manages. Reported objects include those that are configured and collecting data, the types of objects, object counts for adapters, the metrics that are configured and being collected, super metrics, vRealize Operations Manager generated metrics, the number of applications used, and the number of custom groups.

You can use this report to help determine whether the number of objects in your environment exceeds a supported limit.

Where You Audit the System

1   To audit the objects, metrics, applications, and custom groups in your environment, click **Administration**, and then in the left pane click **History > Audit**.

2   Click the **System Audit** tab.

The objects and their associated counts appear in the report.

Table 10-34. System Audit Actions

Option	Description
Download	Download the system information to a report in PDF or XLS format.

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 10-35. System Component Audit Actions

Option	Description
Download	Display the version information in a new browser window.

User Preferences in vRealize Operations Manager

You can configure the user preferences to determine the vRealize Operations Manager display options, such as the number of metrics and groups to display and whether to synchronize system time with the host machine.

To configure the user preferences, in the menu, click the ⬛ icon, and then click **Preferences**. The user preference settings appear in the dialog box.

Table 10-36. User Preference Settings

Option	Description
Display	Configure how many metrics and root cause groups to display. ■ Color scheme: Set the user interface to display in light or dark colors. ■ Important metrics count to show. Set the number of metrics to display. ■ Root cause groups count to show. Set the number of root cause groups to display. ■ Font. Select the font for reports.
Time	Synchronize the time used for the vRealize Operations Manager instance, and display the updated time when vRealize Operations Manager communicates with the host machine. ■ Browser time. All dates and times displayed in the user interface use the time zone settings of the local browser. ■ Host time. All dates and times displayed in the user interface use the time zone of the host machine. ■ Show update time in the application header. Displays the updated time in the top-level header of the vRealize Operations Manager user interface. The updated timestamp appears to the left of the refresh button. Other features, such as dashboards, use the updated time to display data at specific intervals.
Account	Change the password for the user account.

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 a 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 10-37. 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.
Sender E-mail	The email from which the password recovery email is sent.
User name	User name for the STMP 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.

Reset the vRealize Operations Manager Administrator Password on vApp Clusters

You must reset the password if the admin account password is lost.

When the vRealize Operations Manager password for the built-in admin account is lost, follow these steps to reset it on vApp clusters.

Prerequisites

This procedure requires root account credentials.

- In vRealize Operations Manager vApp deployments, when you log in to the console of the virtual application for the first time, you are forced to set a root password.

- The vRealize Operations Manager console root password can be different than the admin account password that you set when configuring the vRealize Operations Manager primary node.

Procedure

1 Log in to the master node command-line console as `root`.

2 Enter the following command, and follow the prompts.

    $VMWARE_PYTHON_BIN $VCOPS_BASE/../vmware-vcopssuite/utilities/
    sliceConfiguration/bin/vcopsSetAdminPassword.py --reset

Reset the vRealize Operations Manager Administrator Password on Windows Clusters

If the admin account password is lost, you must reset the password.

When the vRealize Operations Manager password for the built-in admin account is lost, follow these steps to reset it on Windows clusters.

Procedure

1 Open the command prompt using the **Run as Administrator** option.

2 Enter the following command, and follow the prompts.

    %VMWARE_PYTHON_BIN% %VCOPS_BASE%\..\vmware-
    vcopssuite\utilities\sliceConfiguration\bin\vcopsSetAdminPassword.py --reset

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 the 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. You can also use wildcard certificates in vRealize Operations Manager.

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 can be encrypted by a pass phrase. The generated certificate can be uploaded using the primary node configuration wizard or the administration interface.

- 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 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 create 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 SSL certificates used for establishing a secure connection from vRealize Operations Manager to other services like VMware Identity Manager, vCenter Server, and vRealize Log Insight.

Procedure

1 Generate a Certificate PEM file for use with vRealize Operations Manager.

a Generate a key pair by running this command:

```
openssl genrsa -out key_filename.key 2048
```

b Use the key to generate a certificate signing request by running this command:

```
openssl req -new -key key_filename.key -out certificate_request.csr
```

c Submit the CSR file to your Certificate Authority (CA) to obtain a signed certificate.

d From your Certificate Authority, download the certificate and the complete issuing chain (one or more certificates). Download them in Base64 format.

e   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 user name 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-----
MIIF1DCCBLygAwIBAgIKFYXYUwAAAAAAGTANBgkqhkiG9w0BAQ0FADBhMRMwEQYK
CZImiZPyLGQBGRYDY29tMRUwEwYKCZImiZPyLGQBGRYFdm13Y3MxGDAWBgoJkiaJ
<snip>
vKStQJNr7z2+pTy92M6FgJz3y+daL+9ddbaMNp9fVXjHBoDLGGaLOvyD+KJ8+xba
aGJfGf9ELXM=
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEA4l5ffX694riI1RmdRLJwL6sOWa+Wf70HRoLtx21kZzbXbUQN
mQhTRiidJ3Ro2gRbj/btSsI+OMUzotz5VRT/yeyoTC5l2uJEapld45RroUDHQwWJ
<snip>
DAN9hQus3832xMkAuVP/jt76dHDYyviyIYbmxzMalX7LZy1MCQVg4hCH0vLsHtLh
M1rOAsz62Eht/iB61AsVCCiN3gLrX7MKsYdxZcRVruGXSIh33ynA
-----END RSA PRIVATE KEY-----
```

```
-----BEGIN CERTIFICATE-----
MIIDnTCCAoWgAwIBAgIQY+j29InmdYNCs2cK1H4kPzANBgkqhkiG9w0BAQ0FADBh
MRMwEQYKCZImiZPyLGQBGRYDY29tMRUwEwYKCZImiZPyLGQBGRYFdm13Y3MxGDAW
<snip>
ukzUuqX7wEhc+QgJWgl41mWZBZ09gfsA9XuXBL0k17IpVHpEgwwrjQz8X68m4I99
dD5Pflf/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/0ba6PV8GtWG2svIpcl/eflwGHgTU3zJxR
gkKh7I3K5tGESn81ipyKTkPbYebh+aBMqPKrNNUEKlr0M9sa3WSc0o3350tCc1ew
5ZkNYZ4BRUVYWm0HogeGhOthRn2fAgMBAAECgYABhPmGN3FSZKPDG6HJlARvTlBH
KAGVnBGHd0MOmMAbghFBnBKXa8LwD1dgGBng1oOakEXTftkIjdB+uwkU5P4aRrO7
vGujUtRyRCU/4fjLBDuxQL/KpQfruAQaof9uWUwh5W9fEeW3g26fzVL8AFZnbXS0
7Z0AL1H3LNcLd5rpOQJBANnI7vFu06bFxVF+kq6ZOJFMx7x3K4VGxgg+PfFEBEPS
UJ2LuDH5/Rc63BaxFzM/q3B3Jhehvgw61mMyxU7QSSUCQQC+VDuW3XEWJjSiU6KD
gEGpCyJ5SBePbLSukljpGidKkDNlkLgbWVytCVkTAmuoAz33kMWfqIiNcqQbUgVV
UnpzAkB7d0CPO0deSsy8kMdTmKXLKf4qSF0x55epYK/5MZhBYuA1ENrR6mmjW8ke
TDNc6IGm9sVvrFBz2n9kKYpWThrJAkEAk5R69DtW0cbkLy5MqEzOHQauP36gDi1L
WMXPvUfzSYTQ5aM2rrY2/1FtSSkqUwfYh9sw8eDbqVpIV4rc6dDfcwJBALiiDPT0
tz86wySJNeOiUkQm36iXVF8AckPKT9TrbC3Ho7nC8OzL7gEllETa4Zc86Z3wpcGF
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-----
MIIFWTCCBEGgAwIBAgIKSJGT5gACAAAwKjANBgkqhkiG9w0BAQUFADBBMRMwEQYK
CZImiZPyLGQBGRYDY29tMRYwFAYKCZImiZPyLGQBGRYGdm13YXJlMRIwEAYDVQQD
EwlWTXdhcmUgQ0EwHhcNMTQwMjA1MTg1OTM2WhcNMTYwMjA1MTg1OTM2WjAmMSQw
```

vRealize Operations Manager Certificates

vRealize Operations Manager includes a central page where you can review authentication certificate contents. Certificates allow the vRealize Operations Manager cluster nodes to authenticate each other.

How the Certificates Page Works

The Certificates page lets you examine certificate contents without the need to open the certificate outside of vRealize Operations Manager.

Where You Find Certificates

In the menu, click **Administration**, and then in the left pane, click **Management > Certificates**.

Certificate Tabs

The certificate tab describes columns of exceptions tabs.

Note The CRL tab is enabled only when you select the **Enable Standard Certificate Validation** under **Global Settings**.

Table 10-38. Certificate Tabs

Tabs	Description
Exceptions	Lists the certificate that is accepted by the vRealize Operations Manager administrator but is not certified by the Certificate Authority (CA).
CRL	A Certificate Revocation List (CRL) is a list of digital certificates that have been revoked by the issuing Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. Click the Add icon to upload the certificates.

Certificate Options

The options include a data grid for examining certificate contents.

Table 10-39. Certificate Options

Option	Description
Certificate Thumbprint	Unique alphanumeric string associated with the certificate
Issued By	Content associated with the issuer of the certificate, such as organization name and location
Issued To	Typically, content associated with the issuer, plus the certificate object Identifier (OID)
Expires	The date after which the certificate cannot be used for successful authentication

Add a Custom Certificate to vRealize Operations Manager

If you did not add your own SSL/TLS certificate when configuring the vRealize Operations Manager primary node, you can still add a certificate after vRealize Operations Manager is installed.

Prerequisites

- Create and configure the primary node.

- Verify that your certificate file meets the requirements for vRealize Operations Manager. See the *vRealize Operations Manager vApp Deployment and Configuration Guide* or *vRealize Operations Manager Installation and Configuration Guide for Linux and Windows*.

Procedure

1 In a Web browser, navigate to the vRealize Operations Manager administration interface at https://*node-FQDN-or-ip-address*/admin.

2 Log in with the admin user name and password.

3 At the upper right, click the SSL certificate icon.

4 In the certificate window, click **Install New Certificate**.

5   Click **Browse for certificate**.

6   Locate the certificate `.pem` file, and click **Open** to load the file in the Certificate Information text box.

7   Click **Install**.

Removing an Adapter Certificate

If you want to delete an old or expired certificate associated with an adapter, perform the following steps:

Procedure

1   In a Web browser, navigate to the vRealize Operations Manager administration interface at https://node-FQDN-or-ip-address/ui.

2   Log in with the administrator user name and password.

3   In the menu, click **Administration**, and in the left pane click **Management > Certificates**.

4   In the certificate window, select the certificate that has to be removed.

5   Click **Delete** to remove the certificate.

6   If the certificate is being used by the adapter, then the following message comes up:
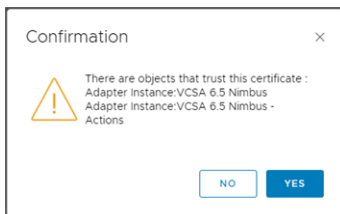
A certificate can be configured for one or more adapters if it is the same destination system.

7   If you delete a certificate which is already being used by another adapter, the adapter fails to connect or start. As a workaround, perform the following steps:

a   On the left pane, click **Solutions**.

b   Select the particular adapter and click the Configure button 🔧 on the toolbar.

c   Click **Test Connection**.

d   A prompt comes up asking the user to import the associated certificate. Click **OK**.

e   Restart the adapter from the **Solutions** page.

Upgrade Internal Certificates

The internal certificates for vRealize Operations Manager expire five years after its initial installation. Upgrade your internal certificates for vRealize Operations Manager 6.3 and later versions using the certificate renewal PAK file. After logging in, if you see a message like, `"vRealize Operations Manager internal certificates will expire on mm/dd/yyyy. Please install a new certificate before the expiry date. For more details,`

`see KB 71018`." in the Quick Start page, you must upgrade your internal certificates for vRealize Operations Manager using the certificate renewal PAK file from the vRealize Operations Manager Administrator interface.

Prerequisites

■ Obtain the PAK file for your cluster. See Obtain the Software Update PAK File for details.

   Note   The certificate renewal PAK is a standalone tool used only for the vRealize Operations Manager internal certificate renew.

■ Bring your cluster offline before installing the PAK file to upgrade your internal certificates.

   Note   If your internal certificates have already expired, install the vRealize_Operations_Manager_Enterprise_Certificate_Renewal_PAK manually. For more information, see the following KB article 71018.

Procedure

1   Log into the vRealize Operations Manager administrator interface of your cluster at `https://master-node-FQDN-or-IP-address/admin`.

2   Take the vRealize Operations Manager cluster offline. For more information, see vRealize Operations Manager Cluster Management.

3   Install the vRealize_Operations_Manager_Enterprise_Certificate_Renewal_PAK to upgrade your internal certificates. See, Install a Software Update.

   Note   After the installation is complete, the administrator interface logs you out.

4   Log back into the vRealize Operations Manager administrator interface.

5   Bring the vRealize Operations Manager cluster back online.

   After the cluster goes online, the upgrade is complete.

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** 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. Some of these settings are not editable. Editable global settings have a hidden Edit icon next to their values. To see the icon, point to the global setting.

Table 10-40. 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 on the adapter. This feature is not implemented in some adapters.
		For example, if the retention time is 360 hours 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 page.
		A value of –1 deletes objects immediately.
		You can define the number of hours per object type to retain objects that no longer exist and check for object type overrides. To add individual object types and set up their values, click the **Object Deletion Scheduling** icon. You can also edit or delete these object types.

Table 10-40. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Deletion Scheduling 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.
Generated Reports Retention	Disabled	Number of months to retain generated reports. If disabled, all the generated reports will be retained.
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 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.
External Event Based Active Symptoms	disabled	The number of days to retain the external event-based active symptoms.

Table 10-40. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Maintain Relationship History		You can maintain a history of all the relationships of all the monitored objects in vRealize Operations Manager .
Dynamic Threshold Calculation	enabled	Determines whether to calculate normal levels of threshold violation for all objects. If the setting is disabled, the following area of vRealize Operations Manager does not work or are not displayed: - Alert symptom definitions based on dynamic thresholds will not work - Metric charts that display normal behavior are not present Disable this setting only if you have no alternative options for managing resource constraints for your vRealize Operations Manager system.
Cost Calculation		The host time at which cost calculations are run.
Customer Experience Improvement Program	enabled	Determines whether to participate in the Customer Experience Improvement Program by having vRealize Operations Manager send anonymous usage data to https://vmware.com.
Allow vCenter users to log in to individual vCenters using the vRealize Operations Manager UI		Lists all the individual vCenter Server on the vRealize Operations Manager login page to allow users to use their individual vCenter Server credentials to log into vRealize Operations Manager. vCenter Server users can log in from vCenter Server clients. Enabled by default, this lists all the configured vCenter Servers in the vRealize Operations Manager login page drop-down.
Allow vCenter users to log in from vCenter clients	enabled	Allows vCenter Server users to log in from the vCenter Server clients.
Allow vCenter users to log in to all vCenters using the vRealize Operations Manager UI	enabled	Lista all the vCenter Servers on the vRealize Operations Manager login page to allow users to use their vCenter Server credentials to log into vRealize Operations Manager. Allows vCenter Server users to log in vRealize Operations Manager UI using any vCenter Server credentials.
System access URL		You can specify the URL that is used to access the system when a load balancer is used. The URL that you enter here is displayed in the outbound notifications and while sharing dashboards.
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 10-40. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Enable Standard Certification Validation		This option enables certificate verification to Test Connection in the Create or Modify AI screen, using a standard verification flow. The option checks CA authority. ■ Certificate Subject DN ■ Subject alternative name ■ Certificate validity period ■ Revocation list 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.
Concurrent UI login sessions	enabled	Allows concurrent UI login sessions per user. Once changed, this setting affects the subsequent login sessions.
Allow non-imported vIDM user access	enabled	Allows non-imported VMware Identity Manager users to be created automatically as read-only users upon first access. If disabled, only VMware Identity Manager imported users or users belonging to imported VMware Identity Manager groups will be granted access.
Currency		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.

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 setting you want to edit.

Note Editable global settings have a hidden **Edit** icon next to their values. To see the icon, point to the global setting.

Table 10-41. Global Settings Options

Option	Description
Edit Global Settings	Click the global setting you want to edit to activate the edit mode and modify the setting values. To edit non-switchable settings, select a value and then click **Save**. To edit switchable settings, select a value and then click **Enable** or **Disable** to change the setting. Click **Cancel** to discard all changes and exit the edit mode.
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.

The Customer Experience Improvement Program

This product participates in VMware's Customer Experience Improvement Program (CEIP). The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. You can choose to join or leave the CEIP for vRealize Operations Manager at any time.

This product also uses a JavaScript operated by VMware's service provider Pendo.io. The JavaScript collects information on your interactions with the user interface such as clickstream data, page loads, limited browser, and device information. It helps VMware to understand how the product is used. This data is used to improve VMware products and services and design them better. For more information, see VMware's Privacy Notices.

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at http://www.vmware.com/trustvmware/ceip.html .

You can opt-out from such data collection by opting out of VMware's CEIP program. Additional controls are also provided to individual users in the user interface.

Join or Leave the Customer Experience Improvement Program for vRealize Operations Manager

You can join or leave the Customer Experience Improvement Program (CEIP) for vRealize Operations Manager at any time.

vRealize Operations Manager gives you the opportunity to join the Customer Experience Improvement Program (CEIP) when you initially install and configure the product. After installation, you can join or leave the CEIP by following these steps.

Procedure

1 In the menu, click **Administration**, and then in the left pane, click **Management > Global Settings**.

2 From the toolbar, click the **Edit** icon.

3 Select or clear the **Customer Experience Improvement Program** option.

   This option activates the program and sends data to www.vmware.com.

4 Click **OK**.

Managing Content

As a vRealize Operations Manager admin, you can take regular backups of your custom and the out-of-the-box content to manage your operational or regulatory needs. If there is a wrong edit or if the need to recover data arises, then you can use the recent backup to restore the content or import the content to a different setup. By taking regular backups, you can also upgrade the vRealize Operations Manager to the latest build without losing or overriding the custom content.

Note Any user with the "View Content Export/Import Page" permission can export the content. However, only a super admin has the privilege to export all the content, including the content owned by other users, for example, custom dashboards.

Creating a Backup

You can create regular backups of your custom and the out-of-the-box content in vRealize Operations Manager . You can use this backup to restore your content or export the content while setting up another environment.

You can take a backup of the following content types available in vRealize Operations Manager .

- Alert Definitions

- Custom Compliance Benchmarks

- Custom Groups

- Dashboards

- Metric Configurations

- Notification Rules

- Policies

- Recommendations

- Report Templates

- Super Metrics

- Symptom Definitions

- Views

Procedure

1 In the **Administration** page, click **Management > Content Management**.

2 In the **Export Content** tab, click **Generate Export Content** to create a backup.

The system compresses the content into one ZIP file.

3 Click the **Download ZIP file** link to download the backup content.

You can use the downloaded content to restore your content or export it to a different setup.

Importing Content

You can take regular backups of your custom and the out-of-the-box content and import it to a different environment.

Prerequisites

- Ensure that you have downloaded the backup ZIP file. For details, see Creating a Backup.

- Ensure that all the users who own the custom dashboards are present in the destination setup so that the custom dashboards are assigned to the respective owners when the content is imported. Otherwise, the dashboards of the owners who are not present in the destination setup will be skipped while importing the content.

Procedure

1 In the **Administration** page, click **Management > Content Management**.

2 Click the **Import Content** tab and then, click **Browse** to select the downloaded ZIP file.

3 If there is a conflict while importing the content, you can select to either **Override exisitng content** or **Skip item(s)**.

The details of the overridden or the skipped content are displayed only during the import and right after the import is completed. You can view this information under the **Results** section in the same page.

4 Click **Import Content**.

After the import is completed, the content is available in the destination setup.

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 10-42. 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.

vRealize Operations Manager Logs for Product UI

How vRealize Operations Manager Logs Work

For troubleshooting in the product UI, the product provides an expandable tree of vRealize Operations Manager log files that you can browse and load for review. You can also edit the log file folders, limit the retained log size, and set logging levels.

vRealize Operations Manager logs are categorized by cluster node, and log type. All logs are in the UTC formatted date and time. The logging format is as follows:

```
Date/Time+0000, LEVEL, [THREAD/IP Address], [Specific Fields], CLASS - MESSAGE
```

If you have configured a timezone for the vRealize Operations Manager VM, the system logs will be in that timezone. The vRealize Operations Manager logs will remain in UTC.

Where You Find vRealize Operations Manager Logs

In the menu, click **Administration**, and in the left pane click **Support > Logs**.

Log Viewer Options

Use the toolbar options to control the tree of items and the viewer.

1   Click **Node** and select any component that is listed under the node.

2   Click the gear icon, enter the logging levels and log size.

3   Click **OK**.

Note   Not all components have relevant syslog information. Therefore, not all nodes have the configuration option enabled.

Figure 10-1. Logs
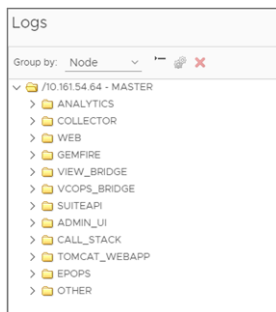
Figure 10-2. Log Options
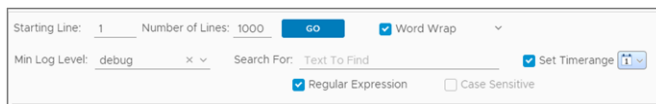
Table 10-43. Log Viewer Toolbar Options

Option	Description
Group By	Organizes the tree by cluster node or log type.
Collapse All	Closes the view of the tree to show only the high-level folders.
Edit Properties	For the selected folder, you can limit the log size and set logging levels.
Delete Selected File	Deletes the log file.
Starting Line	Indicates the starting line of the file .
	0 is for the first line. -1 or no value indicates that the file has to be displayed from the end.

Table 10-43. Log Viewer Toolbar Options (continued)

Option	Description
Number of Lines	Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0.
Min Log Level	If you specify the minimum log level, the logs for that particular log level and higher are shown. For example: If you select **warning**, the logs having the same log level (**warning**) and higher are shown .
Text to Find	Enter the specific text that you want to search in the logs. Add the following filters for search, if required: ■ **Case Sensitive** ■ **Regular Expression** You can perform the search at various levels: ■ On a single file: Use this option if you want to search a single log file . ■ On all the log files of an entity: Use this option if you want to search all the log files of an entity such as a log type or folder. ■ On all the log files of a node: Use this option if you want to search all the log files that are grouped under a node. The last modified time for any file is found by placing the pointer on the file in the tree.
Set Timerange	If you specify a time range, the logs for that particular time range are shown in the search results.
Word Wrap	If you select this option, the part of the line that does not fit on the screen is moved to the next line. If you do not select this option, a scroll bar is provided to see the complete line.

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 > Supoort 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.

vRealize Operations Manager Support Bundles

vRealize Operations Manager support bundles contain log and configuration files that help troubleshoot a vRealize Operations Manager issue.

How Support Bundles Work

Support bundles require that you select nodes or the entire cluster, and the level of logging that you want to collect. After vRealize Operations Manager creates the support bundle, you download it in ZIP format for analysis.

Where You Find Support Bundles

In the menu, click **Administration**, and then in the left pane, select **Support > Support Bundles**.

Support Bundle Options

The options include toolbar and data grid options.

You can click **Add** or click the **Hortizontal Ellipses** to delete, download, or reload support bundles.

Table 10-44. Support Bundle Toolbar Options

Option	Description
Add	Open a dialog box that guides you through the process of creating a support bundle.
Delete	Remove the selected support bundle.
Download	Download the support bundle in ZIP format.
Reload Support Bundles	Refresh the list of support bundles.

Use the data grid options to view item details.

Table 10-45. Support Bundle Data Grid Options

Option	Description
Bundle	System-generated identifier for the support bundle.
Bundle Type	- Light. Include 24 hours of logs. - Full. Include all available logs and configuration files.
Date and Time Created	Time when support bundle creation began.
Status	Progress of support bundle creation.

vRealize Operations Manager Dynamic Thresholds

A threshold marks the boundary between normal and abnormal behavior for a metric. In addition to fixed thresholds, vRealize Operations Manager supports dynamic thresholds for a metric, calculated based on historical and incoming data.

How Dynamic Thresholds Work

By default, dynamic thresholds are refreshed on a regular schedule, but you can recalculate dynamic thresholds outside of the schedule if you want to capture the most recent data.

Where You Find Dynamic Thresholds

In the menu, click **Administration**, and then in the left pane, select **Support > Dynamic Thresholds**.

Dynamic Threshold Options

The dynamic threshold feature includes options to start or stop the calculation process and to review associated values.

Table 10-46. Dynamic Threshold Options

Option	Description
Start	Run the dynamic threshold calculation process now, outside of its normal schedule.
Stop	Stop the dynamic threshold calculation currently in progress.
Calculation progress	Percentage completion of the current dynamic threshold calculation.
Calculation times and Count	Timestamps and metric counts associated with the last dynamic threshold calculation, and the time for the next scheduled calculation.

vRealize Operations Manager Adapter Redescribe

When vRealize Operations Manager redescribes an adapter, vRealize Operations Manager finds the adapter files, gathers information about the abilities of the adapter, and updates the user interface with information about the adapter.

How Adapter Redescribe Works

After installing or updating an adapter, capture the adapter information by having vRealize Operations Manager redescribe its adapters.

Where You Find Adapter Redescribe

In the menu, click **Administration**, and then in left pane, click **Support > Redescribe**.

Adapter Redescribe Options

The feature includes an option to start the adapter describe process.

Table 10-47. Adapter Redescribe Options

Option	Description
Redescribe	Start the adapter describe process.

vRealize Operations Manager provides adapter-specific details from the redescribe process.

Table 10-48. Adapter Redescribe Details

Option	Description
Name	Adapter to which the redescribe process applies.
Status	Success, failure, or other condition related to the last redescribe process.
Describe Version	Version of `describe.xml` against which the last redescribe process ran.
Adapter Version	Version of the adapter against which the last redescribe process ran.
Message	Additional details about the last redescribe process.

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.

Object Type Icons Tab

vRealize Operations Manager obtains data from different sources. Data sources are classified by the type of object or object type. In UI locations where metric data appears for objects, vRealize Operations Manager includes an icon to show the object type. To graphically distinguish the different types of objects, you can customize the icon.

Where You Customize Object Type Icons

In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons > Object Type Icons**.

Table 10-49. Object Type Icons Options

Option	Description
Adapter Type	Icons for all adapters are listed by default. To list a subset of the object types that are valid for one type of adapter, select the adapter type.
Toolbar options	Manages the selected icon. ■ **Upload** uploads a PNG file to uniquely identify the object type. ■ **Assign Default icons** returns the selection to the original icon.
Search	Search for objects with a particular name to narrow the selection of object types displayed.
Object Type	Name of the type of object.
Icon	Pictorial representation of the type of object.

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 the **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.

Adapter Type Icons Tab

Adapters collect and provide data to vRealize Operations Manager. Adapters are classified by the type of adapter or adapter kind. To graphically distinguish the different types of adapters, you can customize the icon.

Where You Customize Adapter Type Icons

In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons > Adapter Type Icons**.

Table 10-50. Adapter Type Icons Options

Option	Description
Toolbar options	Manages the selected icon. ■ **Upload** uploads a PNG file to uniquely identify the adapter type. ■ **Assign Default icons** returns the selection to the original icon.
Name	Name of the type of adapter.
Icon	Pictorial representation of the type of adapter.

Allocate More Virtual Memory to vRealize Operations Manager

You might need to add virtual memory to keep the vRealize Operations Manager process running.

When the vRealize Operations Manager virtual machine requests more memory than is available, the Linux kernel might kill the `vcops-analytics` process, and the product might become unresponsive. If that happens, use the reservation feature in vSphere to specify the guaranteed minimum memory allocation for vRealize Operations Manager virtual machines.

Procedure

1 In the vSphere Client inventory, right-click the vRealize Operations Manager virtual machine and select **Edit Settings**.

2 Click the **Resources** tab, and select **Memory**.

3 Use the **Reservation** option to allocate more memory.

About the vRealize Operations Manager Administration Interface

11

The vRealize Operations Manager administration interface provides access to selected maintenance functions beyond what the product interface supports.

Use the vRealize Operations Manager administration interface instead of the product interface under the following conditions. You can access the administration interface login page from any node in the vRealize Operations Manager analytics cluster by appending **/admin** to the node IP address or FQDN when you enter the URL in your browser.

- Enable or disable high availability (HA).

- Upload and install vRealize Operations Manager software update PAK files.

- The product interface is inaccessible, and you must correct the problem by bringing nodes online, or by restarting nodes or the cluster.

- vRealize Operations Manager needs to be restarted for any reason.

There is some overlap between the administration interface and product interface in terms of access to logs, support bundles, and some of the node maintenance activities that do not involve restarting the cluster, such as adding nodes.

This chapter includes the following topics:

- vRealize Operations Manager Cluster Management

- vRealize Operations Manager Logs for Admin UI

- vRealize Operations Manager Support Bundles

- Update the Reference Database for vRealize Operations Manager

vRealize Operations Manager Cluster Management

vRealize Operations Manager includes a central page where you can monitor and manage the nodes in your vRealize Operations Manager cluster and the adapters that are installed on the nodes.

How Cluster Management Works

You can view and change the online or offline state of the overall vRealize Operations Manager cluster or the individual nodes. In addition, you can enable or disable high availability (HA) and view statistics related to the adapters that are installed on the nodes.

Where You Find Cluster Management

Log in to the vRealize Operations Manager administration interface at https://*master-node-name-or-ip-address*/admin.

Cluster Management Options

The options include cluster-level monitoring and management features.

Table 11-1. Initial Setup Status Details

Option	Description
Cluster Status	Displays the online, offline, or unknown state of the vRealize Operations Manager cluster and provides an option to take the cluster online or offline.
	If a cluster fails to go offline, click the **Force Take Offline** button to take the cluster offline.
	Note The Force Take Offline button appears only when the Bring Cluster offline operation fails.
	You can select to display the reason for taking the cluster offline. Select the **Show reason on maintenance page** check box in the **Take Cluster Offline** dialog box. When you log in to vRealize Operations Manager when the cluster is offline, the reason for taking the cluster offline is displayed.
High Availability	Indicates whether HA is enabled, disabled, or degraded and provides an option to change that setting.
Continuous Availability	Indicates whether CA is enabled, disabled, or degraded and provides an option to change that setting.

vRealize Operations Manager provides node-level information as well as a toolbar for taking nodes online or offline.

Table 11-2. Nodes in the vRealize Operations Manager Cluster

Option	Description
Generate Passphrase	Generate a passphrase that can be used instead of the administrator credentials to add a node to this cluster.
Take Node Online/Offline	You can select the required node and bring it online or offline. You are required to understand the risk involved and provide a valid reason for the action performed when you bring a node online or offline.

Table 11-2. Nodes in the vRealize Operations Manager Cluster (continued)

Option	Description
Reload Nodes	You can fetch data from the nodes.
Shrink Cluster	This option provides a mechanism to remove a node without having to lose any data. The shrink cluster removes nodes by migrating data from one node to any other node.
	All the historical data is either moved to the primary node or any other node, which has sufficient disk space.
	If HA is enabled and you have selected the replica node for removal, then you are asked to select another replica node. vRealize Operations Manager provides a list of nodes that be a possible candidate to become a replica node.
	vRealize Operations Manager stops collecting data from the removed nodes. However, the data that is available in the removed node is migrated to an existing node. Once the migration is complete, then the removed nodes are deleted with the cluster state as offline.
	For remote collectors, if any adapters are on the collectors of the removed nodes, then such nodes are migrated as well.
	Note vRealize Operations Manager cannot move pinned adapters. The adapter instances which were pinned on removed nodes do not move to another collector automatically. You must change the collector before starting the shrink cluster process.

Table 11-3. Nodes in the vRealize Operations Manager Cluster

Option	Description
Node Name	Machine name of the node. The node that you are logged into displays a dot next to the name.
Node Address	Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.
Cluster Role	Type of vRealize Operations Manager node: primary, data, replica, or remote collector.
State	Powered on, powered off, unknown, or other condition of the node.
Status	Online, offline, unknown, or other condition of the node.
Objects	Total environment objects that the node currently monitors.
Metrics	Total metrics that the node has collected since being added to the cluster.

Table 11-3. Nodes in the vRealize Operations Manager Cluster (continued)

Option	Description
Build	vRealize Operations Manager software build number installed on the node.
Version	vRealize Operations Manager software version installed on the node.
Deployment Type	Type of machine on which the node is running: vApp
SSH Status	Enable or disable the SSH Status.

In addition, there are adapter statistics for the selected node.

Table 11-4. Adapters on Server

Option	Description
Name	Name that the installing user gave to the adapter.
Status	Indication of whether the adapter is collecting data or not.
Objects	Total environment objects that the adapter currently monitors.
Metrics	Total metrics that the adapter has collected since being installed on the node.
Last Collection Time	Date and time of the most recent data collection by the adapter.
Added On	Date and time when the adapter was installed on the node.

vRealize Operations Manager Logs for Admin UI

For troubleshooting in the Admin UI, the product provides an expandable tree of vRealize Operations Manager log files that you can browse and load for review.

How vRealize Operations Manager Logs Work

vRealize Operations Manager logs are categorized by cluster node, and functional area or log type.

Where You Find vRealize Operations Manager Logs

1   Log in to the vRealize Operations Manager administration interface at https://*master-node-name-or-ip-address*/admin.

2   In the menu, click **Administration**, and in the left pane click **Support > Logs**.

Log Viewer Options

Use the toolbar options to control the tree of items and the viewer.

Table 11-5. Log Viewer Toolbar Options

Option	Description
Starting Line	Specifies the starting line of the file to be displayed. Note: 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end.
Number of Lines	Specifies the number of lines to be displayed from the file. For example: If you want to see the first 10 lines of the required text, specify the number of lines as 10 and the starting line as 0.
Word Wrap	If you select this option, the extra part of the line that does not fit on the screen is moved to the next line. If you do not select this option, a scroll bar is provided to see the complete line.

vRealize Operations Manager Support Bundles

vRealize Operations Manager support bundles contain log and configuration files that help troubleshoot a vRealize Operations Manager issue.

How Support Bundles Work

Support bundles require that you select nodes or the entire cluster, and the level of logging that you want to collect. After vRealize Operations Manager creates the support bundle, you download it in ZIP format for analysis.

Where You Find Support Bundles

Log in to the vRealize Operations Manager administration interface at https://*master-node-name-or-ip-address*/admin.

Support Bundle Options

The options include toolbar and data grid options.

Use the toolbar options to add, download, or remove items.

Table 11-6. Support Bundle Toolbar Options

Option	Description
Add	Open a dialog box that guides you through the process of creating a support bundle.
Delete	Remove the selected support bundle.
Download	Download the support bundle in ZIP format.
Reload	Refresh the list of support bundles.

Use the data grid options to view item details.

Table 11-7. Support Bundle Data Grid Options

Option	Description
Bundle	System-generated identifier for the support bundle.
Bundle Type	■ Light. Include 24 hours of logs. ■ Full. Include all available logs and configuration files.
Date and Time Created	Time when support bundle creation began.
Status	Progress of support bundle creation.

Update the Reference Database for vRealize Operations Manager

You can update the reference database to have the most updated version of the reference library. The reference database supplies default values for cost calculations.

Procedure

1   In the menu, click **Administration** and in the left pane click **Support > Cost Reference Database**.

    The existing version of the reference database along with the date is displayed.

2   Click **Download Here**.

    The latest version of the reference database is downloaded to the default location.

3   Click **Upload Reference Database** and select the reference database from the default download location.

Results

Note that the updated reference library values are reflected in the cost drivers only after the cost calculation process runs as per the schedule.

OPS-CLI Command-Line Tool

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.

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 report definitions.

- **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 12-1. dashboard Command Options

Command Name	Description	Syntax
dashboard import	Import a dashboard from a file and assign the ownership to a user account.	`dashboard import -h	user-name|all|group:group_name input-file [--force]` ` [--share all|group-name[{,group-name}]] [--retry maxRetryMinutes]` ` [--set rank] [--default] [--create]`
dashboard export	Export an existing dashboard to a file.	`dashboard export -h	user-name dashboard-name [output-dir]`
dashboard defsummary	Import a dashboard from a file and assign the ownership to a user account.	`dashboard defsummary -h	input-file|default` ` --adapterKind adapterKind --resourceKind resourceKind`

Table 12-1. dashboard Command Options (continued)

Command Name	Description	Syntax
dashboard share	Share an existing dashboard with one or multiple user groups.	`dashboard share -h | user-name dahsboard-name all|group-name[{,group-name}]`
dashboard unshare	Stop sharing a dashboard with specified groups.	`dashboard unshare -h | user-name dashboard-name all|group-name[{,group-name}]`
dashboard delete	Permanently delete a dashboard.	`dashboard delete -h | user-name|all|group:group_name dashboard-name`
dashboard reorder	Set the order rank for a dashboard, with an option to make it the default.	`dashboard reorder -h | user-name|all|group:group_name dashboard-name` `                    [--set rank] [--default]`
dashboard show	Show a dashboard.	`dashboard show -h | user-name|all|group:group_name {,dashbaordname}|all`
dashboard hide	Hide a dashboard.	`dashboard hide -h | user-name|all|group:group_name {,dashboardname}|all`

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 12-2. template Command Operations

Command Name	Description	Syntax

template import	Import a template from a file.	`template import -h | input-file`
		`                [--force] [--share all|group-name[{,group-name}]]`
		`                [--retry maxRetryMinutes] [--set rank] [--create]`

template export	Export an existing template to a template file.	`template export -h | template-name [output-dir]`

Table 12-2. template Command Operations (continued)

Command Name	Description	Syntax
template share	Share an existing template with one or multiple user groups.	`template share -h	template-name all|group-name[{,group-name}]`
template unshare	Stop sharing a template with specified groups.	`template unshare -h	template-name all|group-name[{,group-name}]`
template delete	Permanently delete a template.	`template delete -h	template-name`
template reorder	Set the order rank for a template. The order rank controls the order of templates created based on shared templates.	`template reorder -h	template-name [--set rank]`

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 12-3. 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.	`supermetric import -h	input-file` ` [--force] [--policies all|` `policy-name[{,policy-name}]]` ` [--check (true|false)] [--` `retry maxRetryMinutes] [--create]`
supermetric export	Export an existing super metric to a template file.	`supermetric export -h	supermetric-name [output-dir]`

Table 12-3. supermetric Command Operations (continued)

Command Name	Description	Syntax
supermetic configures	Configure properties of a super metric in one or more super metrics packages.	```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}]```
supermetric delete	Permanently delete a super metric.	```supermetric delete -h	supermetric-name```

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 report definitions.

The `report` command uses the following syntax.

```
report -h | import|export|delete [parameters]
```

Table 12-4. report Command Options

Command Name	Description	Syntax
report import	Import a report definition from a file.	`report import -h	input-file [--force]`
report export	Export one or more report definitions to a file.	`report export -h	all|report-name[{,report-name}] [output-dir]`
report delete	Permanently delete one or more report definitions.	`report delete -h	all|report-name[{,report-name}]`

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 12-5. view Command Operations

Command Name	Description	Syntax
view import	Import a view definition from a file.	`view import -h	input-file [--force]`
view export	Export one or more view definitions to a file.	`view export -h	all|view-name[{,view-name}] [output-dir]`
view delete	Permanently delete one or more view definitions.	`view delete -h	all|view-name[{,view-name}]`

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 12-6. file Command Operations

Command Name	Description	Syntax
file import	Import a metric or widget from a file.	```file import -h	reskndmetric|textwidget|topowidget input-file [--title title] [--force]```
file export	Export one or more metrics or text widgets, or export the topology widget to a file.	```file export -h	reskndmetric|textwidget|topowidget all|title[{,title}] [output-dir]```
file delete	Permanently delete a metric or a widget.	```file delete -h	reskndmetric|textwidget|topowidget all|title[{,title}]```
file list	List all metric or a widget files.	```file list -h	reskndmetric|textwidget|topowidget```