vRealize Operations Manager 6.6 Help

22 FEB 2019

vRealize Operations Manager 6.6

vmware®

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

https://docs.vmware.com/

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

Copyright © 2020 VMware, Inc. All rights reserved. Copyright and trademark information.

Contents

VMware vRealize Operations Manager 6.6 Help 8

1 About VMware vRealize Operations Manager 9

2 Planning 10

Reference Architecture 10

Best Practices for Deploying vRealize Operations Manager 10

Initial Considerations for Deploying vRealize Operations Manager 11

Scalability Considerations 13

High Availability Considerations 14

Adapter and Management Packs Considerations 15

Hardware Requirements for Analytic Nodes and Remote Collectors 16

Port Requirements for vRealize Operations Manager 17

Small Deployment Profile for vRealize Operations Manager 20

Medium Deployment Profile for vRealize Operations Manager 21

Large Deployment Profile for vRealize Operations Manager 23

Extra Large Deployment Profile for vRealize Operations Manager 25

Secure Configuration 29

vRealize Operations Manager Security Posture 29

Secure Deployment of vRealize Operations Manager 30

Secure Configuration of vRealize Operations Manager 31

Network Security and Secure Communication 60

Auditing and Logging on your vRealize Operations Manager System 72

3 Installing 74

About Installing 74

Installation Overview 74

Workflow of vRealize Operations Manager Installation 74

Sizing the Cluster 76

Complexity of Your Environment 77

Cluster Nodes 79

About Remote Collector Nodes 80

About High Availability 81

Preparing for Installation 82

Requirements 82

Installing vRealize Operations Manager 86

Deployment of vRealize Operations Manager 87

Installation Types 89

Resize your Cluster by Adding Nodes 95
Gathering More Data by Adding a Remote Collector Node 96
Adding High Availability 98
Cluster and Node Maintenance 99
Post-Installation Considerations 102
About Logging In 102
Secure the Console 103
Log in to a Remote Console Session 104
About New Installations 104
Updating, Migrating and Restoring 106
Obtain the Software Update PAK File 106
Create a Snapshot as Part of an Update 107
How To Preserve Customized Content 107
Backup and Restore 108
Software Updates 108

4 Configuring 113
Connecting to Data Sources 114
VMware vSphere Solution 114
End Point Operations Management Solution 121
Log Insight 180
Business Management 182
vRealize Automation Solution 183
vSAN 189
Installing Optional Solutions 192
Configuring Alerts and Actions 202
All Alerts 202
Types of Alerts 206
Configuring Alerts 215
Viewing Actions 278
Configuring Policies 286
Policies 287
Operational Policies 294
Types of Policies 305
Using the Monitoring Policy Workspace to Create and Modify Operational Policies 318
Define Monitoring Goals for vRealize Operations Manager Solutions 344
Configuring Compliance 346
Defining Compliance Standards 346
Configuring Super Metrics 353
Create a Super Metric 354
Enhancing Your Super Metrics 356

Exporting and Importing a Super Metric 357
Super Metrics Tab 358
Configuring Objects 364
Object Discovery 364
Configuring Data Display 391
Widgets 392
Dashboards 501
Views 527
Reports 543
Configuring Administration Settings 556
License Keys 556
License Groups 557
Maintenance Schedules 558
Manage Maintenance Schedules 559
Managing Users and Access Control 560
Passwords and Certificates 593
Modifying Global Settings 601
Logs 605
Create a Support Bundle 607
Dynamic Thresholds 609
Adapter Redescribe 610
Customizing Icons 611
Allocate More Virtual Memory 613
About the Administration Interface 613
Cluster Status and Troubleshooting 614
Logs 615
Support Bundles 616
Configuring and Using Workload Balance 617
Configuring Workload Balance 617
Using Workload Balance 621
Workload Balance Page 627
Rebalance Containers Dialog Box 628
Manage Rebalance Schedules Dialog Box 629
Workload Automation Policy Settings 629
Using Rebalance Schedules 630

5 Monitoring Objects in Your Managed Environment 631
What to Do When... 631
User Scenario: A User Calls With a Problem 632
User Scenario: An Alert Arrives in Your Inbox 636
User Scenario: You See Problems as You Monitor the State of Your Objects 646

Monitoring and Responding to Alerts 660
Monitoring Alerts 660
Monitoring and Responding to Problems 665
Evaluating Object Information Using Badge Alerts and the Summary Tab 665
Investigating Object Alerts 678
Evaluating Metric Information 685
Analyzing the Resources in Your Environment 696
Using Troubleshooting Tools to Resolve Problems 716
Creating and Using Object Details 722
Examining Relationships in Your Environment 731
User Scenario: Investigate the Root Cause a Problem Using Troubleshooting Tab Options 734
Running Actions from vRealize Operations Manager 738
Run Actions From Toolbars in vRealize Operations Manager 739
Troubleshoot Actions in vRealize Operations Manager 766
Monitor Recent Task Status 768
Troubleshoot Failed Tasks 773
Viewing Your Inventory 781
Inventory Tab 781

6 Planning the Capacity for Your Managed Environment 782
Right-Sizing Capacity for Stress-Free Demand and Value 785
User Scenario: Planning Capacity for an Increase in Workload 789
Create a Sample Project to Increase Workload Capacity 790
Create a Sample Project to Add a Host and Virtual Machines 791
View the Result of Your Capacity Projects 793
Planning Hardware Projects 793
Create a Project to Plan for Hardware Changes 794
Planning Virtual Machine Projects and Scenarios 795
Create a Virtual Machine Project Using Populated Metrics 796
Create a Sample Project for a New Virtual Machine 797
Create a Sample Project to Simulate Removing a Virtual Machine 798
Projects Tab 799
Projects Name and Description Workspace 800
Projects Scenarios Workspace 801
Custom Profiles 802
Custom Profiles Details and Related Policies 803
Custom Profiles Add and Edit Workspace 804
Custom Datacenters 805
Custom Datacenters List 805
Custom Datacenters Add and Edit Workspace 806

7 Metric, Property, and Alert Definitions 808

Metric Definitions in vRealize Operations Manager 808

Metrics for vCenter Server Components 809

Calculated Metrics 887

Self-Monitoring Metrics for vRealize Operations Manager 894

Metrics for vRealize Automation 921

Metrics for vSAN 922

Metrics for the Operating Systems and Remote Service Monitoring Plug-ins in End Point Operations Management 930

Alert Definitions in vRealize Operations Manager 948

Cluster Compute Resource Alert Definitions 949

Host System Alert Definitions 954

vRealize Automation Alert Definitions 968

vSAN Alerts Definitions 969

Alerts in the vSphere Web Client 977

vSphere Distributed Port Group 977

Virtual Machine Alert Definitions 978

vSphere Distributed Switch Alert Definitions 987

vCenter Server Alert Definitions 989

Datastore Alert Definitions 989

Data Center Alert Definitions 995

Custom Data Center Alert Definitions 996

Property Definitions in vRealize Operations Manager 997

Properties for vCenter Server Components 998

Self-Monitoring Properties for vRealize Operations Manager 1012

Properties for vSAN 1014

VMware vRealize Operations Manager 6.6 Help

This documentation contains information for vRealize Operations Manager administrators, virtual infrastructure administrators, and operations engineers who install, configure, and manage objects in your environment.

You can find guidance on commonly performed management activities such as connecting to data sources, configuring users and object groups, responding to alerts, troubleshooting problems, planning capacity, and customizing how data is collected and displayed.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to http://www.vmware.com/support/pubs.

About VMware vRealize Operations Manager

1

With vRealize Operations Manager enterprise software, you can proactively identify and solve emerging issues with predictive analysis and smart alerts, ensuring optimal performance and availability of system resources - across physical, virtual, and cloud infrastructures.

vRealize Operations Manager gives you complete monitoring capability in one place, across applications, storage, and network devices, with an open and extensible platform supported by third-party management packs. In addition, vRealize Operations Manager increases efficiency by streamlining key processes with preinstalled and customizable policies while retaining full control.

Using data collected from system resources (objects), vRealize Operations Manager identifies issues in any monitored system component, often before the customer notices a problem. vRealize Operations Manager also frequently suggests corrective actions you can take to fix the problem right away. For more challenging problems, vRealize Operations Manager offers rich analytical tools that allow you to review and manipulate object data to reveal hidden issues, investigate complex technical problems, identify trends or drill down to gauge the health of a single object.

Planning

2

You plan your environment with recommendations for deployment and secure baseline for the deployment of vRealize Operations Manager.

This chapter includes the following topics:

- Reference Architecture
- Secure Configuration

Reference Architecture

When planning your environment, consider these recommendations for deployment topology, hardware requirements, and interoperability, and scalability.

Best Practices for Deploying vRealize Operations Manager

Implement all best practices when you deploy a production instance of vRealize Operations Manager.

Analytics Nodes

Analytics nodes consist of master nodes, replica nodes, and data nodes.

- Deploy analytics nodes in the same vSphere Cluster.
- Deploy analytics nodes on storage of the same type.
- Depending on the size and performance requirements for analytics nodes, apply Storage DRS Anti-Affinity rules to ensure that nodes are on separate datastores.
- Set Storage DRS to manual for all vRealize Operations Manager analytics nodes.
- If you deploy analytics nodes into a highly consolidated vSphere cluster, configure resource reservation to ensure optimal performance. Ensure that the virtual CPU to physical CPU ratio is not negatively impacting the performance of analytic nodes by validating CPU ready time and CPU co-stop.

- Analytics nodes have a high number of vCPUs to ensure performance of the analytics computation that occurs on each node. Monitor CPU Ready time and CPU Co-Stop to ensure that analytics nodes are not competing for CPU capacity.

If the sizing guideline provides several configurations for the same number of objects, use the configuration which has less number of nodes. For example, if the number of objects is 120, 000, configure the node size as 4 extra large nodes instead of 12 large nodes.

Management Packs and Adapters

Various management packs and adapters have specific configuration requirements. Ensure that you are familiar with all prerequisites before you install a solution and configure the adapter instance.

Red Hat Enterprise Linux (RHEL) OS Installation

- Always follow the RHEL vendor-supplied product installation documentation when installing the OS.
- Firewall protection must be always turned on and for RHEL applications.

Initial Considerations for Deploying vRealize Operations Manager

For the production instance of vRealize Operations Manager to function optimally, your environment must conform to certain configurations. Review and familiarize yourself with these configurations before you deploy a production instance of vRealize Operations Manager.

Sizing

vRealize Operations Manager supports up to 180,000 monitored resources spread across 6 extra large analytic nodes.

Size your vRealize Operations Manager instance to ensure performance and support. For more information about sizing see the following KB article 2150421.

Environment

Deploy analytic nodes in the same vSphere cluster and use identical or similar hosts and storage. If you cannot deploy analytic nodes in the same vSphere cluster, you must deploy them in the same geographical location. vRealize Operations Manager does not support deploying analytics nodes in multiple geographical locations.

Analytics nodes must be able to communicate with one another at all times. The following vSphere events might disrupt connectivity.

- vMotion
- Storage vMotion
- HA
- DRS

Due to a high level of traffic between analytics nodes, all analytics nodes must be Layer 2 Adjacent. Layer 2 Adjacent means that each node is located on the same VLAN and IP subnet, and that VLAN is not stretched between data centers. Latency between analytics nodes cannot exceed 5 milliseconds, and the bandwidth must be equal to or higher than 1 GB per second. It is recommended that bandwidth be 10 GB per second.

If you deploy analytics nodes in to a highly consolidated vSphere cluster, configure resource reservations. A full analytics node, for example a large analytics node that monitors 10,000 resources, requires one virtual CPU to physical CPU. If you experience performance issues, review the CPU ready and co-stop to determine if the virtual to physical CPU ration the cause of the issues. For more information about how to troubleshoot VM performance and interpret CPU performance metrics, see Troubleshooting a virtual machine that has stopped responding: VMM and Guest CPU usage comparison (1017926).

You can deploy remote collectors behind a firewall. You cannot use NAT between remote collectors and analytics nodes.

Multiple Data Centers

If vRealize Operations Manager is monitoring resources in additional data centers, you must use remote collectors and deploy the remote collectors in the remote data center. You might need to modify the intervals at which the configured adapters on the remote collector collect information depending on latency.

It is recommended that latency between sites is less than 200ms. When latency exceeds 200ms, it is recommended that you monitor collections to validate that they are completing in less than five minutes. If collections are not completed in this time limit, increase the interval to 10 minutes.

Certificates

A valid certificate signed by a trusted Certificate Authority, private or public, is an important component when you configure a production instance of vRealize Operations Manager. Configure a Certificate Authority signed certificate against the system before you configure End Point Operations Management agents.

You must include all analytics, remote collectors, and load balancer DNS names in the Subject Alternative Names field of the certificate.

You can configure End Point Operations Management agents to trust the root or intermediate certificate to avoid having to reconfigure all agents if the certificate on the analytics nodes and remote collectors are modified. For more information about root and intermediate certificates, see Specify the End Point Operations Management Agent Setup Properties.

Adapters

It is recommended that you deploy adapters on remote controllers in the same data center as the analytics cluster for large and extra large deployment profiles. Deploying adapters on remote controllers improves performance by reducing load on the analytics node. As an example, you might decide to deploy an adapter remotely if the total resources on a given

analytics node begin to degrade the node's performance. You would likely deploy the adapter on a large remote collector with the appropriate capacity.

You should also deploy adapters to remote collectors when the number of resources the adapters are monitoring exceeds the capacity of the associated analytics node.

Authentication

You can use the Platform Services Controller for user authentication in vRealize Operations Manager. For more information about deploying a highly-available Platform Services Controller instance, see VMware vCenter Server 6.0 Deployment Guide.

Load Balancer

For more information about load balancer configuration, see the vRealize Operations Manager documentation.

Scalability Considerations

Configure your initial deployment of vRealize Operations Manager based on anticipated usage.

Analytics Nodes

Analytics nodes consist of master nodes, master replica nodes, and data nodes.

For enterprise deployments of vRealize Operations Manager, deploy all nodes as large or extra large deployments, depending on your available resources.

Scaling Vertically by Adding Resources

If you deploy analytics nodes in a configuration other than large, you can reconfigure the vCPU and memory. vRealize Operations Manager supports various node sizes.

Table 2-1. Analytics Nodes Deployment Sizes

Node Size	vCPU	Memory
Extra small	2	8 GB
Small	4	16 GB
Medium	8	32 GB
Large	16	48 GB
Extra large	24	128 GB

Scaling Vertically -by Increasing Storage

You can increase storage independently of vCPU and Memory.

To maintain a supported configuration, data nodes deployed in the cluster must be the same node size.

For more information about increasing storage, see Add Disk Space to a vApp Node. You cannot modify the disks of virtual machines that have a snapshot. You must remove all snapshots before you increase disk size.

Scaling Horizontally (Adding nodes)

vRealize Operations Manager 6.6 supports up to 6 extra large analytic nodes in a cluster.

To maintain a supported configuration, analytics nodes deployed in the cluster must be the same node size.

Remote Collectors

vRealize Operations Manager supports two sizes for remote collectors, standard and large. The maximum number of resources is based on the aggregate resources that are collected for all adapters on the remote collector. In large-scale vRealize Operations Manager monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed. Determine the areas of the environment in which the latency is greater than 20 milliseconds and install a remote collector in those areas.

Table 2-2. Supported Remote Collector Sizes

Collector Size	Resources	End Point Operations Management Agents
Standard	1,500	250
Large	12,000	2,500

For more information about sizing see the following KB article 2093783.

High Availability Considerations

HA creates a replica for the vRealize Operations Manager master node and protects the analytics cluster against the loss of a node.

Cluster Management

Clusters consist of master nodes and master replica nodes.

When you enable High Availability, information is stored on the master nodes and master replica nodes.

If the master nodes or master replica nodes are permanently lost, then you must disable and re-enable high availability to reassign the master roles or master replica roles. This process, which includes a hidden cluster rebalance, can take a long time.

Analytics Nodes

Analytics nodes consist of master nodes, master replica nodes, and data nodes.

Enabling High Availability within vRealize Operations Manager is not a disaster recovery solution. Enabling High Availability duplicates data in the system, and doubles the system's compute and capacity requirements. When you enable high availability, you protect vRealize Operations Manager from data loss in the event that a single node is lost. If two or more nodes are lost, the data loss is permanent.

Deploy all analytics nodes to separate hosts to reduce the chance of data loss in the event that a host fails. You can use DRS anti-affinity rules to ensure that VMs remain on separate hosts.

Adapters

In vRealize Operations Manager 6.1 and later, you can create a collector group. A collector group is a collection of nodes (analytic nodes and remote collectors). You can assign adapters to the collector group, rather than assigning an adapter to a single node.

If the node running the adapter fails, the adapter is automatically moved to another node in the collector group.

Assign all normal adapters to collector groups, and not to individual nodes. Do not deploy hybrid adapters in collector groups. For more information about adapters, see the documentation for the specific adapters.

Adapter and Management Packs Considerations

Adapters and management packs have specific configuration considerations.

Normal Adapters

Normal adapters require one-way communication to the monitored endpoint. Deploy normal adapters into collector groups, which are sized to handle failover.

Following is a sample list of adapters provided by VMware for vRealize Operations Manager. Additional adapters can be found on Solutions Exchange.

- vSphere adapter
- Management Pack for NSX for vSphere
- Management Pack for OpenStack
- Management Pack for Storage Devices
- Management Pack for Log Insight

Hybrid Adapters

Hybrid adapters require two-way communication between the adapter and the monitored endpoint.

You must deploy hybrid adapters to a dedicated remote controller. You should configure only one hybrid adapter type for each remote controller. You cannot configure hybrid adapters as part of a collector group. For example, two vRealize Operations for Published Applications adapters can exist on the same node, and two vRealize Operations for Horizon adapters can exist on the same node, but a vRealize Operations for Published Applications adapter and a vRealize Operations for Horizon adapter cannot exist on the same node.

Several hybrid adapters are available for vRealize Operations Manager.

- vRealize Operations for Horizon adapter

- vRealize Operations for Published Applications adapter
- Management Pack for vRealize Hyperic

End Point Operations Management Adapter

By default, End Point Operations Management adapters are installed on all data nodes. Large analytic nodes can support 2,500 agents and large remote collectors can support 2,000 to 10,000 agents for a single cluster. To reduce ingestion load on the cluster, you can point End Point Operations Management adapters at remote collectors. You should assign the dedicated remote collectors to their own collector group, which helps the End Point Operations Management adapter maintain the state of End Point Operations Management resources if a node in the collector group fails.

To reduce the cost of reconfiguring the system, it is recommended that you install End Point Operations Management agents against a DNS entry specific to End Point Operations Management agents if you plan to scale the system beyond a single node.

Remote Collectors Behind a Load Balancer for End Point Operations Management Agents

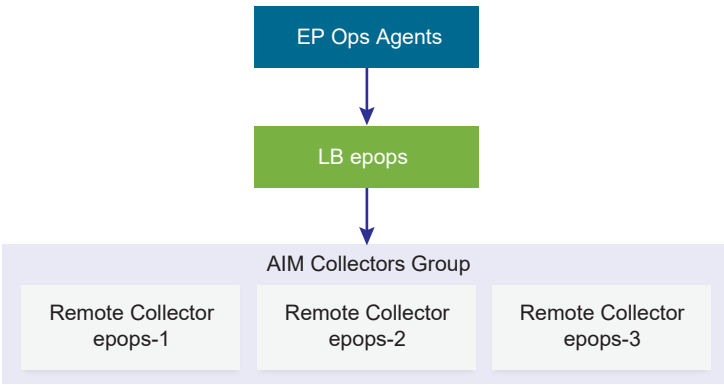

Hardware Requirements for Analytic Nodes and Remote Collectors

Analytics nodes and remote collectors have various hardware requirements for virtual machines and physical machines.

The following table specifies the components to install on each server profile in your deployment, and the required hardware specifications.

Table 2-3. Hardware Requirements for System Components

Server Roles	Virtual CPU	Memory	CPU Requirements	Storage Requirements
Medium analytic node	8 vCPU	32 GB	2.0 Ghz minimum, 2.4 Ghz recommended	1875 IOPS
Large analytic node	16 vCPU	48 GB	2.0 Ghz minimum, 2.4 Ghz recommended	3750 IOPS

Table 2-3. Hardware Requirements for System Components (continued)

Server Roles	Virtual CPU	Memory	CPU Requirements	Storage Requirements
Standard remote collector	2 vCPU	4 GB	2.0 Ghz minimum, 2.4 Ghz recommended	N/A
Large remote collector	4 vCPU	16 GB	2.0 Ghz minimum, 2.4 Ghz recommended	N/A

Storage requirements are based on the maximum supported resources for each node.

vRealize Operations Manager has a high CPU requirement. In general, the more physical CPU that you assign to the analytics cluster, the better the performance. You must use a minimum of eight physical CPU dual socket hosts.

Port Requirements for vRealize Operations Manager

vRealize Operations Manager has certain port requirements for its components. All ports specified are default ports.

Internal Communications

The following components require internal communication.

Table 2-4. Communication Between Master Node and Replica Node

Component	Protocol	Port
Postgres Replica Database	TCP	5433

The XDB ports are required only when you upgrade to vRealize Operations Manager 6.1 or later and are not required for after the upgrade.

Table 2-5. Communication Between Analytics Nodes

Component	Protocol	Port
HTTPS	TCP	443
Gemfire Locator	TCP	6061
Gemfire	TCP	10000
Gemfire	TCP	20000:20010
Cassandra (inter-node)	TCP	7001
Cassandra client	TCP	9042

Table 2-6. Communication From Remote Collector to Analytics Node

Component	Protocol	Port
HTTPS	TCP	443
Gemfire Locator	TCP	6061,
Gemfire	TCP	10000

Table 2-7. Communication Between Remote Collector and Analytics Node

Component	Protocol	Port
HTTPS (Casa)	TCP	443

Table 2-8. Communication Between Remote Collector and Master and Data Nodes

Component	Protocol	Port
HTTP	TCP	80
HTTPS	TCP	443
Gemfire Locator	TCP	6061
Gemfire	TCP and UDP	10000:10010
Gemfire	TCP and UDP	20000:20010
NTP	UDP	123

Table 2-9. Communication From End Point Operations Management Agent to Analytics Node

Component	Protocol	Port
HTTPS	TCP	443

Table 2-10. Communication From End Point Operations Management Agent to Remote Collector

Component	Protocol	Port
HTTPS	TCP	443

External Communications

The following components require external communications.

Table 2-11. Communication from Analytics Nodes and Remote Collectors to External Resources

Component	Protocol	Port
Platform Services Controller	TCP	443
DNS	TCP, UDP	53
LDAP	TCP	389
LDAPS	TCP	636
GC TCP	TCP	3268, 3269
NTP	UDP	123
SMTP	TCP	25

Table 2-11. Communication from Analytics Nodes and Remote Collectors to External Resources (continued)

Component	Protocol	Port
SNMP	UDP	161
Adapters	TCP	**
SSH	TCP	22

** Ports required for adapters to communicate with external devices vary based upon the requirements of the device. Consult adapter documentation for required ports.

Note vROPS requires a TCP connection over HTTP via Port 10433 to connect to vSphere 5.x when retrieving inventory tag information.

Note The user interface and administrative interface to vROPS Operations Manager are through Port 443 with a TCP connection. See additional vROPS port information in the VMware vRealize Operations Manager 6.3 Information Center. Search on "How vRealize Operations Manager Uses Network Ports."

Port Requirements for vRealize Operations Manager

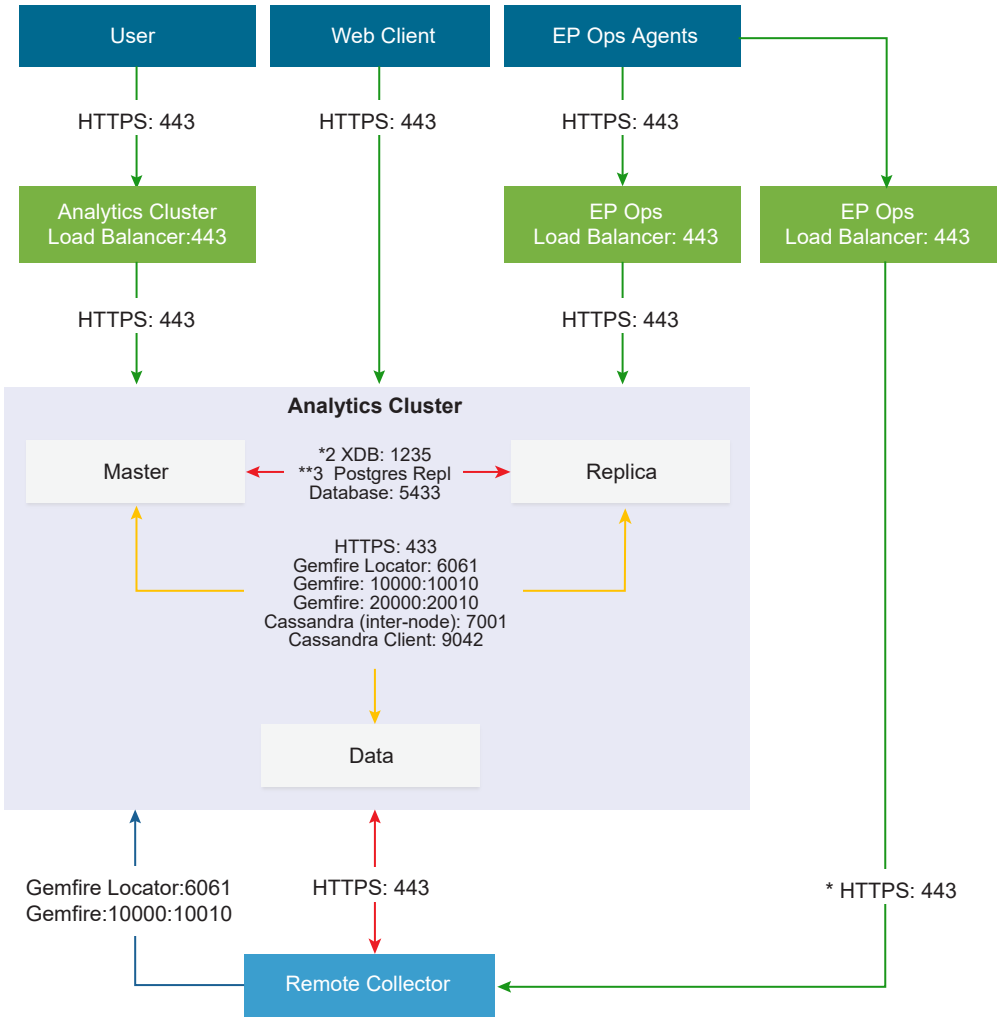

Protocols are not in the diagram.
* Required for upgrading from vRealize Operations Manager 6.0 to 6.1.The ports are closed after the upgrade.
** Required only for High Availability.

Small Deployment Profile for vRealize Operations Manager

The small deployment profile is intended for systems that manage up to 12,000 resources.

Virtual Appliance Name

The small deployment profile contains a single large analytics node, analytic-1.ra.local.

Deployment Profile Support

The small deployment profile supports the following configuration.

- 12,000 resources
- 1,000 End Point Operations Management agents
- Data retention for six months

Additional DNS Entries

You can add additional DNS entries for your organization's future requirements. If you do not expect your planned deployment to exceed a single node, you can configure End Point Operations Management agents against the analytics nodes.

epops.ra.local -> analytic-1.ra.local

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *epops.refarch.local*
- DNS Name = *analytic-1.ra.local*

This is an example of a small deployment profile.

Table 2-12. Adapter Properties

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytic-1	A	2,000
DEFAULT	analytic-1	B	4,000
DEFAULT	analytic-1	C	2,000
DEFAULT	analytic-1	D	3,000

vRealize Operations Manager Small Deployment Profile Architecture

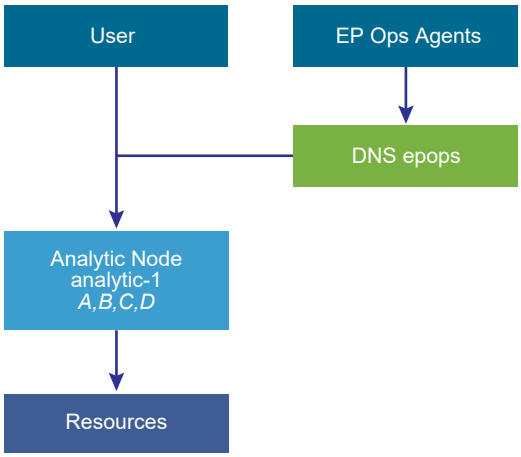

Medium Deployment Profile for vRealize Operations Manager

The medium deployment profile is intended for systems that manage 40,000 resources, 20,000 of which are enabled for High Availability. In the medium deployment profile, adapters are deployed on the analytics nodes by default. If you experience problems with data ingestion, move these adapters to remote controllers.

Virtual Appliance Names

The medium deployment profile contains eight medium analytics nodes.

- analytic-1.ra.lcoal
- analytic-2.ra.lcoal
- analytic-3.ra.lcoal
- analytic-4.ra.lcoal
- analytic-5.ra.lcoal
- analytic-6.ra.lcoal
- analytic-7.ra.lcoal
- analytic-8.ra.lcoal

Deployment Profile Support

The medium deployment profile supports the following configuration.

- 40,000 total resources, 20,000 enabled for HA
- 6,000 End Point Operations Management agents
- Data retention for six months

Load Balanced Addresses

- analytics.ra.local
- epops.ra.local

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *epops.refarch.local*
- DNS Name = *analytic-1.ra.local*

This is an example of a medium deployment profile.

Table 2-13. Adapter Properties

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytic-1	A	2,000
DEFAULT	analytic-2	B	4,000
DEFAULT	analytic-3	C	2,000
DEFAULT	analytic-4	D	3,000
DEFAULT	analytic-5	E	1,000
DEFAULT	analytic-6	F	2,000

Table 2-13. Adapter Properties (continued)

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytic-7	G	1,500
DEFAULT	analytic-8	H	4,500

vRealize Operations Manager Medium Deployment Profile Architecture

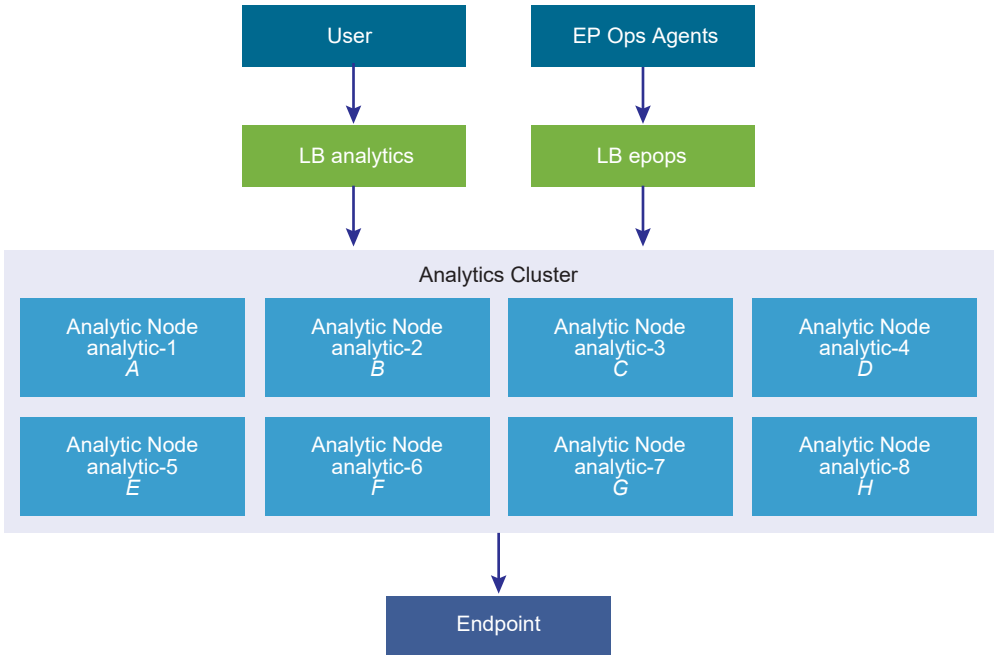

Large Deployment Profile for vRealize Operations Manager

The large deployment profile is intended for systems that manage 80,000 resources, 40,000 of which are enabled with High Availability. All adapters are deployed to remote controllers in large deployment profiles to offload CPU usage from the analytics cluster.

Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, large remote collectors for adapters, and large remote collectors for End Point Operations Management agents.

- analytic-1.ra.lcoal
- analytic-2.ra.lcoal
- analytic-3.ra.lcoal
- analytic-4.ra.lcoal
- analytic-5.ra.lcoal
- analytic-6.ra.lcoal
- analytic-7.ra.lcoal

- analytic-8.ra.lcoal

Deployment Profile Support

The large deployment profile supports the following configuration.

- 80,000 total resources, 40,000 enabled for HA
- 10,000 End Point Operations Management agents
- Data retention for six months

Load Balanced Addresses

- `analytics.ra.local`
- `epops.ra.local`

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytic.refarch.local*
- DNS Name = *epops.refarch.local*
- DNS Name = *analytic-1.ra.local* to DNS Name = *analytic-8.ra.local*
- DNS Name = *remote-1.ra.local* to DNS Name = *remote-N.ra.local*
- DNS Name = *epops-1.ra.lcoal* to DNS Name = *epops-N.ra.local*

This is an example of a large deployment profile.

Table 2-14. Adapter Properties

Collector Group	Remote Collector	Adapter	Resources	End Point Operations Management Agents
1	remote-1	A	5,000	N/A
1	remote-2	B	5,000	N/A
		Total	10,000	N/A
2	remote-3	C	10,000	N/A
2	remote-4	D	5,000	N/A
2	remote-5	E	5,000	N/A
		Total	20,000	N/A
AIM	epops-1	epops	4,800	800
	epops-2	epops	4,800	800
		Total	9,600	1,600

If a remote collector is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit of 10,000 resource for each remote collector.

The estimate of 9,600 resources uses six resources for each End Point Operations Management agent.

vRealize Operations Manager Large Deployment Profile Architecture

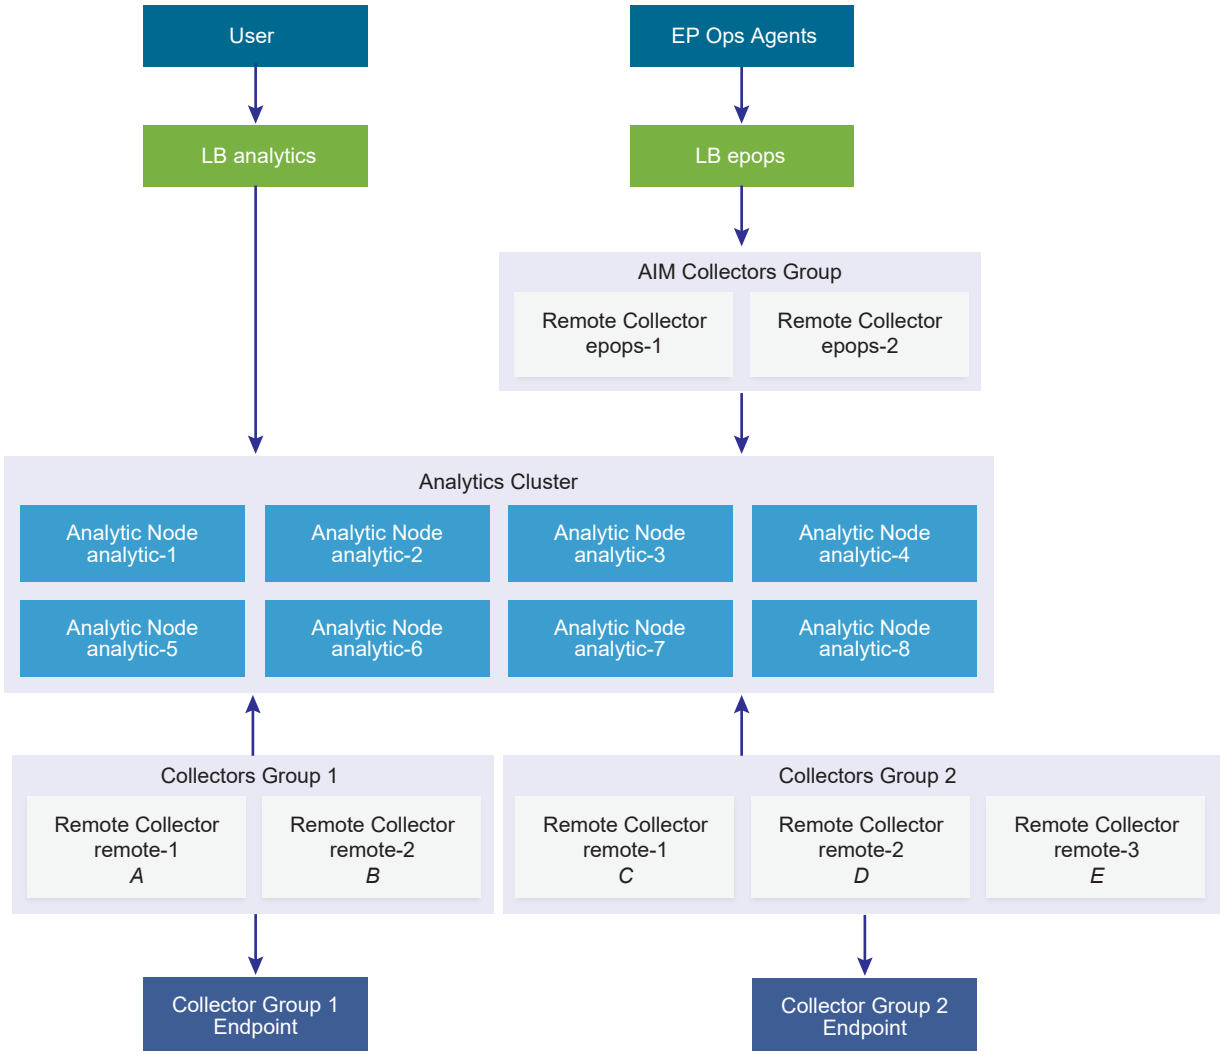

Extra Large Deployment Profile for vRealize Operations Manager

The extra large deployment profile is intended for systems that manage 180,000 resources, 90,000 of which are enabled for High Availability. This deployment is divided into two data centers and is the maximum supported analytics cluster deployment.

Virtual Appliance Names

The extra large deployment profile contains 6 extra large analytics nodes, X large remote collectors for adapters, and Y large remote collectors for End Point Operations Management agents.

- `analytic-1.ra.local`
- `analytic-2.ra.local`

- analytic-3.ra.local
- analytic-4.ra.local
- analytic-5.ra.local
- analytic-6.ra.local

Deployment Profile Support

- 180,000 total resources, 90,000 enabled for HA
- 10,000 End Point Operations Management agents
- Data retention for six months

Load Balanced Addresses

- analytics.ra.local
- epops-a.ra.local
- epops-b.ra.local

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytic.refarch.local*
- DNS Name = *epops-a.refarch.local*
- DNS Name = *epops-b.refarch.local*
- DNS Name = *analytic-1.ra.local* to *analytic-16.ra.local*
- DNS Name = *remote-1.ra.local* to *remote-N.ra.local*
- DNS Name = *epops-1.ra.local* to *epops-N.ra.local*

This is an example of an extra large deployment profile. The adapter in the example provides N-1 redundancy, meaning, if two adapters support 20,000 resources, then a third adapter is added to attain a supported configuration that allows for a single failure.

Table 2-15. Adapter Properties

Collector Group	Data Center	Remote Collector	Adapter	Resources	End Point Operations Management agents
1	A	remote-1	A	5,000	N/A
1	A	remote-2	B	5,000	N/A
			Total	10,000	
2	A	remote-3	C	2,000	N/A
2	A	remote-3	D	2,000	N/A

Table 2-15. Adapter Properties (continued)

Collector Group	Data Center	Remote Collector	Adapter	Resources	End Point Operations Management agents
2	A	remote-3	E	1,000	N/A
2	A	remote-4	F	7,000	N/A
2	A	remote-5	G	8,000	N/A
2	A	remote-6	H	5,000	N/A
2	A	remote-7	I	6,000	N/A
			Total	31,000	
3	B	remote-8	J	10,000	N/A
3	B	remote-9	K	5,000	N/A
3	B	remote-10	L	5,000	N/A
			Total	20,000	
AIM-1	A	epops-1	epops	8,004	1,334
AIM-1	A	epops-2	epops	7,998	1,333
	A	epops-3	epops	7,998	1,333
			Total	24,000	4,000
AIM-2	B	epops-4	epops	8,004	1,334
AIM-2	B	epops-5	epops	7,998	1,333
AIM-2	B	epops-6	epops	7,998	1,333
			Total	24,000	4,000

If a remote collector is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit of 15,000 resource for each remote collector.

The estimate of 24,000 resources for AIM-1 and AIM-2 collector groups uses six resources for each End Point Operations Management agent.

vRealize Operations Manager Extra Large Deployment Profile Architecture - Data Center A

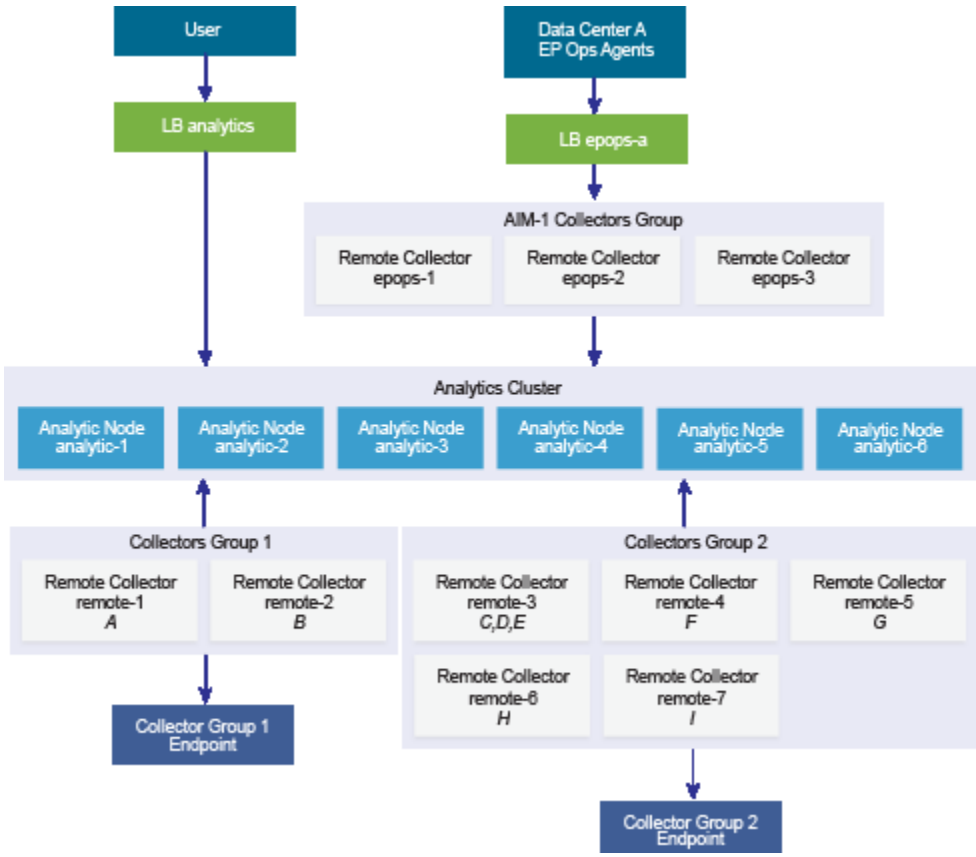

vRealize Operations Manager Extra Large Deployment Profile Architecture - Data Center B

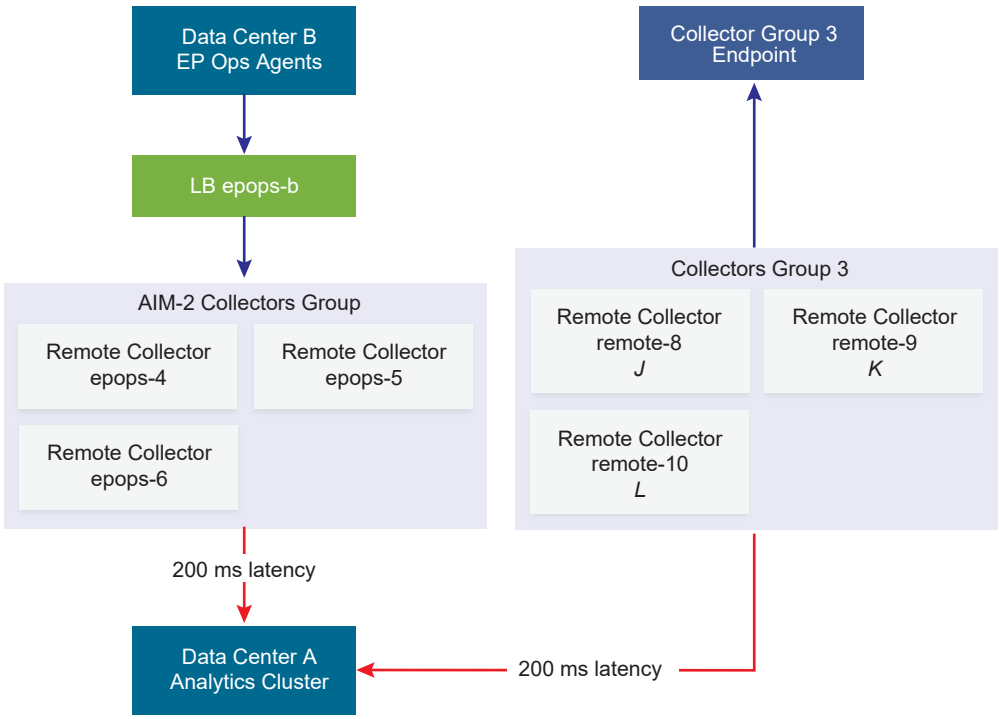

Secure Configuration

Ensure you meet the security requirements in your environment with the recommendations provided.

vRealize Operations Manager Security Posture

The security posture of vRealize Operations Manager assumes a complete secure environment based on system and network configuration, organizational security policies, and best practices. It is important that you perform the hardening activities according to your organization's security policies and best practices.

The document is broken down into the following sections:

- Secure Deployment
- Secure Configuration
- Network Security
- Communication

The guide details the installation of the Virtual Application.

To ensure that your system is securely hardened, review the recommendations and assess them against your organization's security policies and risk exposure.

Secure Deployment of vRealize Operations Manager

You must verify the integrity of the installation media before you install the product to ensure authenticity of the downloaded files.

Verify the Integrity of Installation Media

After you download the media, use the MD5/SHA1 sum value to verify the integrity of the download. Always verify the SHA1 hash after you download an ISO, offline bundle, or patch to ensure the integrity and authenticity of the downloaded files. If you obtain physical media from VMware and the security seal is broken, return the software to VMware for a replacement.

Procedure

- Compare the MD5/SHA1 hash output with the value posted on the VMware Web site.

  SHA1 or MD5 hash should match.

  Note The vRealize Operations Manager `6.x–x.pak` files are signed by the VMware software publishing certificate. vRealize Operations Manager validates the signature of the PAK file before installation.

Hardening the Deployed Software Infrastructure

As part of your hardening process, you must harden the deployed software infrastructure that supports your VMware system.

Before you harden your VMware system, review and address security deficiencies in your supporting software infrastructure to create a completely hardened and secure environment. Software infrastructure elements to consider include operating system components, supporting software, and database software. Address security concerns in these and other components according to the manufacturer's recommendations and other relevant security protocols.

Hardening the VMware vSphere Environment

vRealize Operations Manager relies on a secure VMware vSphere environment to achieve the greatest benefits and a secured infrastructure.

Assess the VMware vSphere environment and verify that the appropriate level of vSphere hardening guidance is enforced and maintained.

For more guidance about hardening, see http://www.vmware.com/security/hardening-guides.html.

Reviewing Installed and Unsupported Software

Vulnerabilities in unused software might increase the risk of unauthorized system access and disruption of availability. Review the software that is installed on VMware host machines and evaluate its use.

Do not install software that is not required for the secure operation of the system on any of the vRealize Operations Manager node hosts. Uninstall unused or nonessential software.

Installing unsupported, untested, or unapproved software on infrastructure products such as vRealize Operations Manager is a threat to the infrastructure.

To minimize the threat to the infrastructure, do not install or use any third-party software that is not supported by VMware on VMware supplied hosts.

Assess your vRealize Operations Manager deployment and inventory of installed products to verify that no unsupported software is installed.

For more information about the support policies for third-party products, see the VMware support at http://www.vmware.com/security/hardening-guides.html.

Verify Third-Party Software

Do not use third-party software that VMware does not support. Verify that all third-party software is securely configured and patched in accordance with third-party vendor guidance.

Inauthentic, insecure, or unpatched vulnerabilities of third-party software installed on VMware host machines might put the system at risk of unauthorized access and disruption of availability. All software that VMware does not supply must be appropriately secured and patched.

If you must use third-party software that VMware does not support, consult the third-party vendor for secure configuration and patching requirements.

VMware Security Advisories and Patches

VMware occasionally releases security advisories for products. Being aware of these advisories can ensure that you have the safest underlying product and that the product is not vulnerable to known threats.

Assess the vRealize Operations Manager installation, patching, and upgrade history and verify that the released VMware Security Advisories are followed and enforced.

It is recommended that you always remain on the most recent vRealize Operations Manager release, as this will include the most recent security fixes also.

For more information about the current VMware security advisories, see http://www.vmware.com/security/advisories/.

Secure Configuration of vRealize Operations Manager

As a security best practice, you must secure the vRealize Operations Manager console and manage Secure Shell (SSH), administrative accounts, and console access. Ensure that your system is deployed with secure transmission channels.

You must also follow certain security best practices for running End Point Operations Management agents.

Secure the vRealize Operations Manager Console

After you install vRealize Operations Manager, you must log in for the first time and secure the console of each node in the cluster.

Prerequisites

Install vRealize Operations Manager.

Procedure

1. Locate the node console in vCenter or by direct access.

   In vCenter, press Alt+F1 to access the login prompt. For security reasons, vRealize Operations Manager remote terminal sessions are disabled by default.

2. Log in as root.

   vRealize Operations Manager does not allow you to access the command prompt until you create a root password.

3. At the password prompt, press **Enter**.
4. At the old password prompt, press **Enter**.
5. At the prompt for a new password, enter the root password that you want and note it for future reference.
6. Reenter the root password.
7. Log out of the console.

Change the Root Password

You can change the root password for any vRealize Operations Manager master or data node at any time by using the console.

The root user bypasses the `pam_cracklib` module password complexity check, which is found in `etc/pam.d/common–password`. All hardened appliances enable `enforce_for_root` for the `pw_history` module, found in the `etc/pam.d/common–password` file. The system remembers the last five passwords by default. Old passwords are stored for each user in the `/etc/security/opasswd` file.

Prerequisites

Verify that the root password for the appliance meets your organization's corporate password complexity requirements. If the account password starts with `$6$`, it uses a sha512 hash. This is the standard hash for all hardened appliances.

Procedure

1. Run the `# passwd` command at the root shell of the appliance.
2. To verify the hash of the root password, log in as root and run the `# more /etc/shadow` command.

   The hash information appears.

3. If the root password does not contain a sha512 hash, run the `passwd` command to change it.

Manage Password Expiry

Configure all account password expirations in accordance with your organization's security policies.

By default, all hardened VMware appliances use a 60-day password expiry. On most hardened appliances, the root account is set to a 365-day password expiry. As a best practice, verify that the expiry on all accounts meets security and operation requirements standards.

If the root password expires, you cannot reinstate it. You must implement site-specific policies to prevent administrative and root passwords from expiring.

Procedure

1. Log in to your virtual appliance machines as root and run the `# more /etc/shadow` command to verify the password expiry on all accounts.
2. To modify the expiry of the root account, run the `# passwd –x 365 root` command.

   In this command, 365 specifies the number of days until password expiry. Use the same command to modify any user, substituting the specific account for `root` and replacing the number of days to meet the expiry standards of the organization.

   By default, the root password is set for 365 days.

Managing Secure Shell, Administrative Accounts, and Console Access

For remote connections, all hardened appliances include the Secure Shell (SSH) protocol. SSH is disabled by default on the hardened appliance.

SSH is an interactive command-line environment that supports remote connections to a vRealize Operations Manager node. SSH requires high-privileged user account credentials. SSH activities generally bypass the role-based access control (RBAC) and audit controls of the vRealize Operations Manager node.

As a best practice, disable SSH in a production environment and enable it only to diagnose or troubleshoot problems that you cannot resolve by other means. Leave it enabled only while needed for a specific purpose and in accordance with your organization's security policies. If you enable SSH, ensure that it is protected against attack and that you enable it only for as long as required. Depending on your vSphere configuration, you can enable or disable SSH when you deploy your Open Virtualization Format (OVF) template.

As a simple test to determine whether SSH is enabled on a machine, try to open a connection by using SSH. If the connection opens and requests credentials, then SSH is enabled and is available for making connections.

Secure Shell Root User

Because VMware appliances do not include preconfigured default user accounts, the root account can use SSH to directly log in by default. Disable SSH as root as soon as possible.

To meet the compliance standards for nonrepudiation, the SSH server on all hardened appliances is preconfigured with the AllowGroups wheel entry to restrict SSH access to the secondary group wheel. For separation of duties, you can modify the AllowGroups wheel entry in the `/etc/ssh/sshd_config` file to use another group such as sshd.

The wheel group is enabled with the `pam_wheel` module for superuser access, so members of the wheel group can use the su-root command, where the root password is required. Group separation enables users to use SSH to the appliance, but not to use the su command to log in as root. Do not remove or modify other entries in the AllowGroups field, which ensures proper appliance function. After making a change, restart the SSH daemon by running the `# service sshd restart` command.

Enable or Disable Secure Shell on a vRealize Operations Manager node

You can enable Secure Shell (SSH) on a vRealize Operations Manager node for troubleshooting. For example, to troubleshoot a server, you might require console access to the server. This is through SSH. Disable SSH on a vRealize Operations Manager node for normal operation.

Procedure

1. Access the console of the vRealize Operations Manager node from vCenter.
2. Press Alt + F1 to access the login prompt then log in.
3. Run the `#chkconfig` command.
4. If the sshd service is off, run the `#chkconfig sshd on` command.
5. Run the `#service sshd start` command to start the sshd service.
6. Run the `#service sshd stop` command to stop the sshd service.

Create a Local Administrative Account for Secure Shell

You must create local administrative accounts that can be used as Secure Shell (SSH) and that are members of the secondary wheel group, or both before you remove the root SSH access.

Before you disable direct root access, test that authorized administrators can access SSH by using `AllowGroups`, and that they can use the wheel group and the `su` command to log in as root.

Procedure

1. Log in as root and run the following commands.

   ```
   # useradd –d /home/vropsuser –g users –G wheel -m
   # passwd username
   ```

   Wheel is the group specified in `AllowGroups` for SSH access. To add multiple secondary groups, use `–G wheel,sshd`.

2 Switch to the user and provide a new password to ensure password complexity checking.

```
# su - username
username@hostname:~>passwd
```

If the password complexity is met, the password updates. If the password complexity is not met, the password reverts to the original password, and you must rerun the password command.

After you create the login accounts to allow SSH remote access and use the `su` command to log in as root using the wheel access, you can remove the root account from the SSH direct login.

3 To remove direct login to SSH, modify the `/etc/ssh/sshd_config` file by replacing `(#)PermitRootLogin yes` with `PermitRootLogin no`.

What to do next

Disable direct logins as root. By default, the hardened appliances allow direct login to root through the console. After you create administrative accounts for nonrepudiation and test them for wheel access (`su–root`), disable direct root logins by editing the `/etc/securetty` file as root and replacing the `tty1` entry with `console`.

Restrict Secure Shell Access

As part of your system hardening process, restrict Secure Shell (SSH) access by configuring the tcp_wrappers package appropriately on all VMware virtual appliance host machines. Also maintain required SSH key file permissions on these appliances.

All VMware virtual appliances include the tcp_wrappers package to allow tcp-supported daemons to control the network subnets that can access the libwrapped daemons. By default, the `/etc/hosts.allow` file contains a generic entry, `sshd: ALL : ALLOW`, that allows all access to the secure shell. Restrict this access as appropriate for your organization.

Procedure

1 Open the `/etc/hosts.allow` file on your virtual appliance host machine in a text editor.

2 Change the generic entry in your production environment to include only the local host entries and the management network subnet for secure operations.

```
sshd:127.0.0.1 : ALLOW
sshd: [::1] : ALLOW
sshd: 10.0.0.0 :ALLOW
```

In this example, all local host connections and connections that the clients make on the 10.0.0.0 subnet are allowed.

3 Add all appropriate machine identification, for example, host name, IP address, fully qualified domain name (FQDN), and loopback.

4 Save the file and close it.

Maintain Secure Shell Key File Permissions

To maintain an appropriate level of security, configure Secure Shell (SSH) key file permissions.

Procedure

1. View the public host key files, located in `/etc/ssh/*key.pub`.
2. Verify that these files are owned by root, that the group is owned by root, and that the files have permissions set to 0644.

   The permissions are (-rw-r--r--).
3. Close all files.
4. View the private host key files, located in `/etc/ssh/*key`.
5. Verify that root owns these files and the group, and that the files have permissions set to 0600.

   The permissions are (-rw-------).
6. Close all files.

Harden the Secure Shell Server Configuration

Where possible, the Virtual Application Installation (OVF) has a default hardened configuration. Users can verify that their configuration is appropriately hardened by examining the server and client service in the global options section of the configuration file.

If possible, restrict use of the SSH server to a management subnet in the `/etc/hosts.allow` file.

Procedure

1. Open the `/etc/ssh/sshd_config` server configuration file and verify that the settings are correct.

Setting	Status
Server Daemon Protocol	Protocol 2
Ciphers	Ciphers aes256-ctr,aes128-ctr
TCP Forwarding	AllowTCPForwarding no
Server Gateway Ports	Gateway Ports no
X11 Forwarding	X11Forwarding no
SSH Service	Use the AllowGroups field and specify a group permitted to access and add members to the secondary group for users permitted to ue the service.
GSSAPI Authentication	GSSAPIAuthentication no, if unused
Kerberos Authentication	KerberosAuthentication no, if unused
Local Variables (AcceptEnv global option)	Set to `disabled by commenting out` or `enabled for only LC_* or LANG variables`
Tunnel Configuration	PermitTunnel no

Setting	Status
Network Sessions	MaxSessions 1
Strict Mode Checking	Strict Modes yes
Privilege Separation	UsePrivilegeSeparation yes
rhosts RSA Authentication	RhostsRSAAuthentication no
Compression	Compression delayed or Compression no
Message Authentication code	MACs hmac-sha1
User Access Restriction	PermitUserEnvironment no

2 Save your changes and close the file.

Harden the Secure Shell Client Configuration

As part of your system hardening monitoring process, verify hardening of the SSH client by examining the SSH client configuration file on virtual appliance host machines to ensure that it is configured according to VMware guidelines.

Procedure

1 Open the SSH client configuration file, `/etc/ssh/ssh_config`, and verify that the settings in the global options section are correct.

Setting	Status
Client Protocol	Protocol 2
Client Gateway Ports	Gateway Ports no
GSSAPI Authentication	GSSAPIAuthentication no
Local Variables (SendEnv global option)	Provide only LC_* or LANG variables
CBC Ciphers	Ciphers aes256-ctr,aes128-ctr
Message Authentication Codes	Used in the `MACs hmac-sha1` entry only

2 Save your changes and close the file.

Disable Direct Logins as Root

By default, the hardened appliances allow you to use the console to log in directly as root. As a security best practice, you can disable direct logins after you create an administrative account for nonrepudiation and test it for wheel access by using the `su-root` command.

Prerequisites

- Complete the steps in the topic called Create a Local Administrative Account for Secure Shell.
- Verify that you have tested accessing the system as an administrator before you disable direct root logins.

Procedure

1 Log in as root and navigate to the `/etc/securetty` file.

   You can access this file from the command prompt.

2 Replace the `tty1` entry with `console`.

Disable SSH Access for the Admin User Account

As a security best practice, you can disable SSH access for the admin user account. The vRealize Operations Manager admin account and the Linux admin account share the same password. Disabling SSH access to the admin user enforces defense in depth by ensuring all users of SSH first login to a lesser privileged service account with a password that differs from the vRealize Operations Manager admin account and then switch user to a higher privilege such as the admin or root.

Procedure

1 Edit the `/etc/ssh/sshd_config` file.

   You can access this file from the command prompt.

2 Add the `DenyUsers admin` entry anywhere in the file and save the file.

3 To restart the sshd server, run the `service sshd restart` command.

Set Boot Loader Authentication

To provide an appropriate level of security, configure boot loader authentication on your VMware virtual appliances. If the system boot loader requires no authentication, users with console access to the system might be able to alter the system boot configuration or boot the system to single user or maintenance mode, which can result in denial of service or unauthorized system access.

Because boot loader authentication is not set by default on the VMware virtual appliances, you must create a GRUB password to configure it.

Procedure

1 Verify whether a boot password exists by locating the `password --md5 <password-hash>` line in the `/boot/grub/menu.lst` file on your virtual appliances.

2 If no password exists, run the `# /usr/sbin/grub-md5-crypt` command on your virtual appliance.

   An MD5 password is generated, and the command supplies the md5 hash output.

3 Append the password to the `menu.lst` file by running the `# password --md5 <hash from grub-md5-crypt>` command.

Single-User or Maintenance Mode Authentication

If the system does not require valid root authentication before it boots into single-user or maintenance mode, anyone who invokes single-user or maintenance mode is granted privileged access to all files on the system.

Procedure

- Review the `/etc/inittab` file and ensure that the following two lines appear: `ls:S:wait:/etc/init.d/rc S` and `~~:S:respawn:/sbin/sulogin`.

Monitor Minimal Necessary User Accounts

You must monitor existing user accounts and ensure that any unnecessary user accounts are removed.

Procedure

- Run the *host*`:~ # cat /etc/passwd` command and verify the minimal necessary user accounts:

```
bin:x:1:1:bin:/bin:/bin/bash
daemon:x:2:2:Daemon:/sbin:/bin/bash
haldaemon:x:101:102:User for haldaemon:/var/run/hald:/bin/false
mail:x:8:12:Mailer daemon:/var/spool/clientmqueue:/bin/false
man:x:13:62:Manual pages viewer:/var/cache/man:/bin/bash
messagebus:x:100:101:User for D-Bus:/var/run/dbus:/bin/false
nobody:x:65534:65533:nobody:/var/lib/nobody:/bin/bash
ntp:x:74:106:NTP daemon:/var/lib/ntp:/bin/false
polkituser:x:103:104:PolicyKit:/var/run/PolicyKit:/bin/false
postfix:x:51:51:Postfix Daemon:/var/spool/postfix:/bin/false
root:x:0:0:root:/root:/bin/bash
sshd:x:71:65:SSH daemon:/var/lib/sshd:/bin/false
suse-ncc:x:104:107:Novell Customer Center User:/var/lib/YaST2/suse-ncc-fakehome:/bin/bash
uuidd:x:102:103:User for uuidd:/var/run/uuidd:/bin/false
wwwrun:x:30:8:WWW daemon apache:/var/lib/wwwrun:/bin/false
nginx:x:105:108:user for nginx:/var/lib/nginx:/bin/false
admin:x:1000:1003::/home/admin:/bin/bash
tcserver:x:1001:1004:tc Server User:/home/tcserver:/bin/bash
postgres:x:1002:100::/var/vmware/vpostgres/9.3:/bin/bash
```

Monitor Minimal Necessary Groups

You must monitor existing groups and members to ensure that any unnecessary groups or group access is removed.

Procedure

- Run the `<host>:~ # cat /etc/group` command to verify the minimum necessary groups and group membership.

```
audio:x:17:
bin:x:1:daemon
cdrom:x:20:
console:x:21:
```

```
daemon:x:2:
dialout:x:16:u1,tcserver,postgres
disk:x:6:
floppy:x:19:
haldaemon:!:102:
kmem:x:9:
mail:x:12:
man:x:62:
messagebus:!:101:
modem:x:43:
nobody:x:65533:
nogroup:x:65534:nobody
ntp:!:106:
polkituser:!:105:
public:x:32:
root:x:0:admin
shadow:x:15:
sshd:!:65:
suse-ncc:!:107:
sys:x:3:
tape:!:103:
trusted:x:42:
tty:x:5:
utmp:x:22:
uuidd:!:104:
video:x:33:u1,tcserver,postgres
wheel:x:10:root,admin
www:x:8:
xok:x:41:
maildrop:!:1001:
postfix:!:51:
users:x:100:
vami:!:1002:root
nginx:!:108:
admin:!:1003:
vfabric:!:1004:admin,wwwrun
```

Resetting the vRealize Operations Manager Administrator Password (Linux)

As a security best practice, you can reset the vRealize Operations Manager password on Linux clusters for vApp or Linux installations.

Procedure

1. Log in to the remote console of the master node as root.
2. Enter the `$VMWARE_PYTHON_BIN $VCOPS_BASE/../vmware-vcopssuite/utilities/sliceConfiguration/bin/vcopsSetAdminPassword.py --reset` command and follow the prompts.

Configure NTP on VMware Appliances

For critical time sourcing, disable host time synchronization and use the Network Time Protocol (NTP) on VMware appliances. You must configure a trusted remote NTP server for time

synchronization. The NTP server must be an authoritative time server or at least synchronized with an authoritative time server.

The NTP daemon on VMware virtual appliances provides synchronized time services. NTP is disabled by default, so you need to configure it manually. If possible, use NTP in production environments to track user actions and to detect potential malicious attacks and intrusions through accurate audit and log keeping. For information about NTP security notices, see the NTP Web site.

The NTP configuration file is located in the `/etc/ntp.conf` file on each appliance.

Procedure

1. Navigate to the `/etc/ntp.conf` configuration file on your virtual appliance host machine.
2. Set the file ownership to **`root:root`**.
3. Set the permissions to **`0640`**.
4. To mitigate the risk of a denial-of-service amplification attack on the NTP service, open the `/etc/ntp.conf` file and ensure that the restrict lines appear in the file.

   ```
   restrict default kod nomodify notrap nopeer noquery
   restrict -6 default kod nomodify notrap nopeer noquery
   restrict 127.0.0.1
   restrict -6 ::1
   ```

5. Save any changes and close the files.

   For information on NTP security notices, see http://support.ntp.org/bin/view/Main/SecurityNotice.

Disable the TCP Timestamp Response on Linux

Use the TCP timestamp response to approximate the remote host's uptime and aid in further attacks. Additionally, some operating systems can be fingerprinted based on the behavior of their TCP time stamps.

Procedure

- Disable the TCP timestamp response on Linux.
  a. To set the value of `net.ipv4.tcp_timestamps` to `0`, run the `sysctl -w net.ipv4.tcp_timestamps=0` command.
  b. Add the `ipv4.tcp_timestamps=0` value in the default `sysctl.conf` file.

Enable FIPS 140-2 Mode

The version of OpenSSL that is shipped with vRealize Operations Manager 6.3 and later releases is FIPS 140-2 certified. However, the FIPS mode is not enabled by default.

You can enable the FIPS mode if there is a security compliance requirement to use FIPS certified cryptographic algorithms with the FIPS mode enabled.

Procedure

1 To replace the `mod_ssl.so` file run the following command:

```
cd /usr/lib64/apache2-prefork/
cp mod_ssl.so mod_ssl.so.old
cp mod_ssl.so.FIPSON.openssl1.0.2 mod_ssl.so
```

2 Modify your Apache2 configuration by editing the `/etc/apache2/ssl-global.conf` file.

3 Search for the `<IfModule mod_ssl.c>` line and add the `SSLFIPS on` directive below it.

4 To reset the Apache configuration, run the `service apache2 restart` command.

TLS for Data in Transit

As a security best practice, ensure that the system is deployed with secure transmission channels.

Configure Strong Protocols for vRealize Operations Manager

Protocols such as SSLv2 and SSLv3 are no longer considered secure. In addition, it is recommended that you disable TLS 1.0. Enable only TLS 1.1 and TLS 1.2.

Verify the Correct Use of Protocols in Apache HTTPD

vRealize Operations Manager disables SSLv2 and SSLv3 by default. You must disable weak protocols on all load balancers before you put the system into production.

Procedure

1 Run the `grep SSLProtocol /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf | grep -v '#'` command from the command prompt to verify that SSLv2 and SSLv3 are disabled.

   If the protocols are disabled, the command returns the following output: `SSLProtocol All -SSLv2 -SSLv3`

2 To also disable the TLS 1.0 protocol, run the `sed -i "/^[^#]*SSLProtocol/ c\SSLProtocol All -SSLv2 -SSLv3 -TLSv1" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` command from the command prompt.

3 To restart the Apache2 server, run the `/etc/init.d/apache2 restart` command from the command prompt.

Verify the Correct Use of Protocols in the GemFire TLS Handler

vRealize Operations Manager disables SSLv3 by default. You must disable weak protocols on all load balancers before you put the system into production.

Procedure

1 Verify that the protocols are enabled. To verify that the protocols are enabled, run the following commands on each node:

```
grep cluster-ssl-protocol /usr/lib/vmware-vcops/user/conf/gemfire.properties | grep -v '#'
```

The following result is expected:

```
cluster-ssl-protocols=TLSv1.2 TLSv1.1 TLSv1
grep cluster-ssl-protocol /usr/lib/vmware-vcops/user/conf/gemfire.native.properties | grep -v '#'
```

The following result is expected:

```
cluster-ssl-protocols=TLSv1.2 TLSv1.1 TLSv1
grep cluster-ssl-protocol /usr/lib/vmware-vcops/user/conf/gemfire.locator.properties | grep -v '#'
```

The following result is expected:

```
cluster-ssl-protocols=TLSv1.2 TLSv1.1 TLSv1
```

2 Disable TLS 1.0.

a Navigate to the administrator user interface at `url/admin` .

b Click **Bring Offline**.

c To disable SSLv3 and TLS 1.0, run the following commands:

```
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2
TLSv1.1" /usr/lib/vmware-vcops/user/conf/gemfire.properties
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2
TLSv1.1" /usr/lib/vmware-vcops/user/conf/gemfire.native.properties
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2
TLSv1.1" /usr/lib/vmware-vcops/user/conf/gemfire.locator.properties
```

Repeat this step for each node

d Navigate to the administrator user interface.

e Click **Bring Online**.

3 Reenable TLS 1.0.

a Navigate to the administrator user interface to bring the cluster offline: `url/admin`.

b Click **Bring Offline**.

c To ensure that SSLv3 and TLS 1.0 are disabled, run the following commands:

```
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2 TLSv1.1
TLSv1" /usr/lib/vmware-vcops/user/conf/gemfire.properties
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2 TLSv1.1
TLSv1" /usr/lib/vmware-vcops/user/conf/gemfire.native.properties
sed -i "/^[^#]*cluster-ssl-protocol/ c\cluster-ssl-protocols=TLSv1.2 TLSv1.1
TLSv1" /usr/lib/vmware-vcops/user/conf/gemfire.locator.properties
```

Repeat this step for each node.

d Navigate to the administrator user interface to bring the cluster online.

e Click **Bring Online**.

Configure vRealize Operations Manager to Use Strong Ciphers

For maximum security, you must configure vRealize Operations Manager components to use strong ciphers. To ensure that only strong ciphers are selected, disable the use of weak ciphers. Configure the server to support only strong ciphers and to use sufficiently large key sizes. Also, configure the ciphers in a suitable order.

vRealize Operations Manager disables the use of cipher suites using the DHE key exchange by default. Ensure that you disable the same weak cipher suites on all load balancers before you put the system into production.

Using Strong Ciphers

The encryption cipher negotiated between the server and the browser determines the key exchange method and encryption strength that is used in a TLS session.

Verify the Correct Use of Cipher Suites in Apache HTTPD

For maximum security, verify the correct use of cipher suites in Apache httpd.

Procedure

1. To verify the correct use of cipher suites in Apache httpd, run the `grep SSLCipherSuite /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf | grep -v '#'` command from the command prompt.

   If Apache httpd uses the correct cipher suites, the command returns the following output: `SSLCipherSuite kECDH+AESGCM:ECDH+AESGCM:RSA+AESGCM:kECDH+AES:ECDH+AES:RSA+AES:!aNULL!ADH:!EXP:!MD5:!3DES:!CAMELLIA:!PSK:!SRP:!DH`

2. To configure the correct use of cipher suites, run the `sed -i "/^[^#]*SSLCipherSuite/ c \SSLCipherSuite kECDH+AESGCM:ECDH+AESGCM:RSA+AESGCM:kECDH+AES:ECDH+AES:RSA+AES:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP:\!DH" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` command from the command prompt.

   Run this command if the output in Step 1 is not as expected.

   This command disables all cipher suites that use DH and DHE key exchange methods.

3. Run the `/etc/init.d/apache2 restart` command from the command prompt to restart the Apache2 server.

4. To reenable DH, remove `!DH` from the cipher suites by running the `sed -i "/^[^#]*SSLCipherSuite/ c\SSLCipherSuite kECDH+AESGCM:ECDH+AESGCM:RSA+AESGCM:kECDH+AES:ECDH+AES:RSA+AES:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` command from the command prompt.

5. Run the `/etc/init.d/apache2 restart` command from the command prompt to restart the Apache2 server.

Verify the Correct Use of Cipher Suites in GemFire TLS Handler

For maximum security, verify the correct use of cipher suites in GemFire TLS Handler.

Procedure

1 To verify that the cipher suites are enabled, run the following commands on each node to verify that the protocols are enabled:

```
grep cluster-ssl-ciphers /usr/lib/vmware-vcops/user/conf/gemfire.properties |
grep -v '#'
```

```
grep cluster-ssl-ciphers /usr/lib/vmware-vcops/user/conf/
gemfire.native.properties | grep -v '#'
```

```
grep cluster-ssl-ciphers /usr/lib/vmware-vcops/user/conf/
gemfire.locator.properties | grep -v '#'
```

2 Configure the correct cipher suites.

   a Navigate to the administrator user interface at *URL*/admin.

   b To bring the cluster offline, click **Bring Offline**.

   c To configure the correct cipher suites, run the following commands:

```
sed -i "/^[^#]*cluster-ssl-ciphers/ c\cluster-ssl-
ciphers=TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256" /usr/lib/vmware-vcops/user/
conf/gemfire.properties
```

```
sed -i "/^[^#]*cluster-ssl-ciphers/ c\cluster-ssl-
ciphers=TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256" /usr/lib/vmware-vcops/user/
conf/gemfire.native.properties
```

```
sed -i "/^[^#]*cluster-ssl-ciphers/ c\cluster-ssl-
ciphers=TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256" /usr/lib/vmware-vcops/user/
conf/gemfire.locator.properties
```

   Repeat this step for each node.

   d Navigate to the administrator user interface at *URL*/admin.

   e Click **Bring Online**.

Enabling TLS on Localhost Connections

By default, the localhost connections to the PostgreSQL database do not use TLS. To enable TLS, you have to either generate a self-signed certificate with OpenSSL or provide your own certificate.

To enable TLS on localhost connections to PostgreSQL, complete the following steps:

1 Generate or Provide Your Own Self-Signed Certificate with OpenSSL

2 Install the Certificate for PostgreSQL

3 Enable TLS on PostgreSQL

Generate or Provide Your Own Self-Signed Certificate with OpenSSL

Localhost connections to the PostgreSQL database do not use TLS. To enable TLS, you can generate your own self-signed certificate with OpenSSL or provide your own certificate.

- To generate a self-signed certificate with OpenSSL, run the following commands:

```
openssl req –new –text –out cert.req
 openssl rsa –in privkey.pem –out cert.pem
 openssl req –x509 –in cert.req –text –key cert.pem –out cert.cert
```

- To provide your own certificate, complete the following steps:
  - Modify the ownership of the `CAcerts.crt` file to `postgres`.
  - Edit the `postgresql.conf` file to include the directive `ssl_ca_file = 'CAcerts.crt`.

    If you are using a certificate with a CA chain, you must add a `CAcerts.crt` file containing the intermediate and root CA certificates to the same directory.

Install the Certificate for PostgreSQL

You must install the certificate for PostgreSQL when you enable TLS on localhost connections to PostgreSQL.

Procedure

1. Copy the `cert.pem` file to `/storage/db/vcops/vpostgres/data/server.key`.
2. Copy the `cert.cert` file to `/storage/db/vcops/vpostgres/data/server.crt`.
3. Run the `chmod 600 /storage/db/vcops/vpostgres/data/server.key` command.
4. Run the `chmod 600 /storage/db/vcops/vpostgres/data/server.crt` command.
5. Run the `chown postgres /storage/db/vcops/vpostgres/data/server.key` and `chown postgres /storage/db/vcops/vpostgres/data/server.crt` commands to change the ownership of the `server.crt` and `server.key` files from `root` to `postgres`.

Enable TLS on PostgreSQL

You must edit the `postgresql.conf` file to enable TLS on localhost connections to PostgreSQL.

Procedure

- Edit the `postgresql.conf` file at `/storage/db/vcops/vpostgres/data/` and make the following changes:
  a. Set `ssl = on`.
  b. Set `ssl_cert_file = 'server.crt'`.
  c. Set `ssl_key_file = 'server.key'`.

Application Resources That Must be Protected

As a security best practice, ensure that the application resources are protected.

Follow the steps to ensure that the application resources are protected.

Procedure

1 Run the `Find / –path /proc –prune –o –type f –perm +6000 –ls` command to verify that the files have a well defined SUID and GUID bits set.

The following list appears:

```
354131   24 –rwsr–xr–x   1 polkituser root 23176 /usr/lib/PolicyKit/polkit–set–default–helper
354126   20 –rwxr–sr–x   1 root     polkituser    19208 /usr/lib/PolicyKit/polkit–grant–helper
354125   20 –rwxr–sr–x   1 root     polkituser    19008 /usr/lib/PolicyKit/polkit–explicit–grant–
helper
354130   24 –rwxr–sr–x   1 root     polkituser    23160 /usr/lib/PolicyKit/polkit–revoke–helper
354127   12 –rwsr–x–––   1 root     polkituser    10744 /usr/lib/PolicyKit/polkit–grant–helper–pam
354128   16 –rwxr–sr–x   1 root     polkituser    14856 /usr/lib/PolicyKit/polkit–read–auth–helper
 73886   84 –rwsr–xr–x   1 root     shadow      77848 /usr/bin/chsh
 73888   88 –rwsr–xr–x   1 root     shadow      85952 /usr/bin/gpasswd
 73887   20 –rwsr–xr–x   1 root     shadow      19320 /usr/bin/expiry
 73890   84 –rwsr–xr–x   1 root     root        81856 /usr/bin/passwd
 73799  240 –rwsr–xr–x   1 root     root       238488 /usr/bin/sudo
 73889   20 –rwsr–xr–x   1 root     root        19416 /usr/bin/newgrp
 73884   92 –rwsr–xr–x   1 root     shadow      86200 /usr/bin/chage
 73885   88 –rwsr–xr–x   1 root     shadow      82472 /usr/bin/chfn
 73916   40 –rwsr–x–––   1 root     trusted     40432 /usr/bin/crontab
296275   28 –rwsr–xr–x   1 root     root        26945 /usr/lib64/pt_chown
353804  816 –r–xr–sr–x   1 root     mail       829672 /usr/sbin/sendmail
278545   36 –rwsr–xr–x   1 root     root        35792 /bin/ping6
278585   40 –rwsr–xr–x   1 root     root        40016 /bin/su
278544   40 –rwsr–xr–x   1 root     root        40048 /bin/ping
278638   72 –rwsr–xr–x   1 root     root        69240 /bin/umount
278637  100 –rwsr–xr–x   1 root     root        94808 /bin/mount
475333   48 –rwsr–x–––   1 root     messagebus    47912 /lib64/dbus–1/dbus–daemon–launch–helper
 41001   36 –rwsr–xr–x   1 root     shadow      35688 /sbin/unix_chkpwd
 41118   12 –rwsr–xr–x   1 root     shadow      10736 /sbin/unix2_chkpwd
```

2 Run the `find / –path */proc –prune –o –nouser –o –nogroup` command to verify that all the files in the vApp have an owner.

All the files have an owner if there are no results.

3 Run the `find / –name "*.*" –type f –perm –a+w | xargs ls –ldb` command to verify that none of the files are world writable files by reviewing permissions of all the files on the vApp.

None of the files must include the permission `xx2`.

4 Run the `find / –path */proc –prune –o ! –user root –o –user admin –print` command to verify that the files are owned by the correct user.

All the files belong to either `root` or `admin` if there are no results.

5 Run the `find /usr/lib/vmware–casa/ –type f –perm –o=w` command to ensure that files in the `/usr/lib/vmware–casa/` directory are not world writable.

There must be no results.

6 Run the `find /usr/lib/vmware-vcops/ -type f -perm -o=w` command to ensure that files in the `/usr/lib/vmware-vcops/` directory are not world writable.

   There must be no results.

7 Run the `find /usr/lib/vmware-vcopssuite/ -type f -perm -o=w` command to ensure that files in the `/usr/lib/vmware-vcopssuite/` directory are not world writable.

   There must be no results.

Configure PostgreSQL Client Authentication

You can configure the system for client authentication. You can configure the system for local trust authentication. This allows any local user, including the database super user to connect as a PostgreSQL user without a password. If you want to provide a strong defense and if you do not have significant trust in all local user accounts, use another authentication method. The md5 method is set by default. Verify that md5 is set for all local and host connections.

You can find the client authentication configuration settings for the postgres service instance in `/storage/db/vcops/vpostgres/data/pg_hba.conf`. Verify that md5 is set for all local and host connections.

The client authentication configuration settings for the postgres-repl service instance can be found in `/storage/db/vcops/vpostgres/repl/pg_hba.conf`. Verify that md5 is set for all local and host connections.

Note Do not modify client configuration settings for the postgres user account.

Apache Configuration

Disable Web Directory Browsing

As a security best practice, ensure that a user cannot bowse through a directory because it can increase the risk of exposure to directory traversal attacks.

Procedure

- Verify that web directory browsing is disabled for all directories.
  a Open the `/etc/apache2/default-server.conf` and `/usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` files in a text editor.
  b Verify that for each `<Directory>` listing, the option called `Indexes` for the relevant tag is omitted from the `Options` line.

Remove the Sample Code for the Apache2 Server

Apache includes two sample Common Gateway Interface (CGI) scripts, `printenv` and `test-cgi`. A production Web server must contain only components that are operationally necessary. These components have the potential to disclose critical information about the system to an attacker.

As a security best practice, delete the CGI scripts from the `cgi-bin` directory.

Procedure

- To remove `test-cgi` and `prinenv` scripts, run the `rm /usr/share/doc/packages/apache2/test-cgi` and `rm /usr/share/doc/packages/apache2/printenv` commands.

Verify Server Tokens for the Apache2 Server

As part of your system hardening process, verify server tokens for the Apache2 server. The Web server response header of an HTTP response can contain several fields of information. Information includes the requested HTML page, the Web server type and version, the operating system and version, and ports associated with the Web server. This information provides malicious users important information without the use of extensive tools.

The directive `ServerTokens` must be set to `Prod`. For example, `ServerTokens Prod`. This directive controls whether the response header field of the server that is sent back to clients includes a description of the operating system and information about compiled-in modules.

Procedure

1. To verify server tokens, run the `cat /etc/apache2/sysconfig.d/global.conf | grep ServerTokens` command.
2. To modify `ServerTokens OS` to `ServerTokens Prod`, run the `sed -i 's/\(ServerTokens\s\+\)OS/\1Prod/g' /etc/apache2/sysconfig.d/global.conf` command.

Disable the Trace Method for the Apache2 Server

In standard production operations, use of diagnostics can reveal undiscovered vulnerabilities that lead to compromised data. To prevent misuse of data, disable the HTTP `Trace` method.

Procedure

1. To verify the `Trace` method for the Apache2 server, run the following command `grep TraceEnable /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf`.
2. To disable the `Trace` method for the Apache2 server, run the following command `sed -i "/^[^#]*TraceEnable/ c\TraceEnable off" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf`.

Disable Configuration Modes

As a best practice, when you install, configure, or maintain vRealize Operations Manager, you can modify the configuration or settings to enable troubleshooting and debugging of your installation.

Catalog and audit each of the changes you make to ensure that they are properly secured. Do not put the changes into production if you are not sure that your configuration changes are correctly secured.

Managing Nonessential Software Components

To minimize security risks, remove or configure nonessential software from your vRealize Operations Manager host machines.

Configure all software that you do not remove in accordance with manufacturer recommendations and security best practices to minimize its potential to create security breaches.

Secure the USB Mass Storage Handler

Secure the USB mass storage handler to prevent it from loading by default on vRealize appliances and to prevent its use as the USB device handler with the vRealize appliances. Potential attackers can exploit this handler to install malicious software.

Procedure

1. Open the`/etc/modprobe.conf.local` file in a text editor.
2. Ensure that the `install usb-storage /bin/true` line appears in the file.
3. Save the file and close it.

Secure the Bluetooth Protocol Handler

Secure the Bluetooth protocol handler on your vRealize Appliances to prevent potential attackers from exploiting it.

Binding the Bluetooth protocol to the network stack is unnecessary and can increase the attack surface of the host. Prevent the Bluetooth protocol handler module from loading by default on vRealize Appliances.

Procedure

1. Open the `/etc/modprobe.conf.local` file in a text editor.
2. Ensure that the line `install bluetooth /bin/true` appears in this file.
3. Save the file and close it.

Secure the Stream Control Transmission Protocol

Prevent the Stream Control Transmission Protocol (SCTP) module from loading on vRealize appliances by default. Potential attackers could exploit this protocol to compromise your system.

Configure your system to prevent the SCTP module from loading unless it is absolutely necessary. SCTP is an unused IETF-standardized transport layer protocol. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1. Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the following line appears in this file.

```
install sctp /bin/true
```

3 Save the file and close it.

Secure the Datagram Congestion Control Protocol

As part of your system hardening activities, prevent the Datagram Congestion Control Protocol (DCCP) module from loading on vRealize appliances by default. Potential attackers can exploit this protocol to compromise your system.

Avoid loading the DCCP module, unless it is absolutely necessary. DCCP is a proposed transport layer protocol, which is not used. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes can cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the DCCP lines appear in the file.

```
install dccp /bin/true
install dccp_ipv4 /bin/true
install dccp_ipv6 /bin/true
```

3 Save the file and close it.

Secure Reliable Datagram Sockets Protocol

As part of your system hardening activities, prevent the Reliable Datagram Sockets (RDS) protocol from loading on your vRealize appliances by default. Potential attackers can exploit this protocol to compromise your system.

Binding the RDS protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the `install rds /bin/true` line appears in this file.

3 Save the file and close it.

Secure the Transparent Inter-Process Communication Protocol

As part of your system hardening activities, prevent the Transparent Inter-Process Communication protocol (TIPC) from loading on your virtual appliance host machines by default. Potential attackers can exploit this protocol to compromise your system.

Binding the TIPC protocol to the network stack increases the attack surface of the host. Unprivileged local processes can cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the `install tipc /bin/true` line appears in this file.

3 Save the file and close it.

Secure Internet Packet Exchange Protocol

Prevent the Internetwork Packet Exchange (IPX) protocol from loading vRealize appliances by default. Potential attackers could exploit this protocol to compromise your system.

Avoid loading the IPX protocol module unless it is absolutely necessary. IPX protocol is an obsolete network-layer protocol. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the system to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the line `install ipx /bin/true` appears in this file.

3 Save the file and close it.

Secure Appletalk Protocol

Prevent the Appletalk protocol from loading on vRealize appliances by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the Appletalk Protocol module unless it is absolutely necessary. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the system to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the line `install appletalk /bin/true` appears in this file.

3 Save the file and close it.

Secure DECnet Protocol

Prevent the DECnet protocol from loading on your system by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the DECnet Protocol module unless it is absolutely necessary. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes could cause the system to dynamically load a protocol handler by using the protocol to open a socket.

Procedure

1 Open the DECnet Protocol `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the line `install decnet /bin/true` appears in this file.

3 Save the file and close it.

Secure Firewire Module

Prevent the Firewire module from loading on vRealize appliances by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the Firewire module unless it is absolutely necessary.

Procedure

1 Open the `/etc/modprobe.conf.local` file in a text editor.

2 Ensure that the line `install ieee1394 /bin/true` appears in this file.

3 Save the file and close it.

Kernel Message Logging

The `kernel.printk` specification in the `/etc/sysctl.conf` file specifies the kernel print logging specifications.

There are 4 values specified:

- `console loglevel`. The lowest priority of messages printed to the console.
- `default loglevel`. The lowest level for messages without a specific log level.
- The lowest possible level for the console log level.
- The default value for console log level.

There are eight possible entries per value.

- `define KERN_EMERG "<0>" /* system is unusable */`
- `define KERN_ALERT "<1>" /* action must be taken immediately */`
- `define KERN_CRIT "<2>" /* critical conditions */`
- `define KERN_ERR "<3>" /* error conditions */`
- `define KERN_WARNING "<4>" /* warning conditions */`
- `define KERN_NOTICE "<5>" /* normal but significant condition */`
- `define KERN_INFO "<6>" /* informational */`
- `define KERN_DEBUG "<7>" /* debug-level messages */`

Set the `kernel.printk` values to **`3 4 1 7`** and ensure that the line `kernel.printk=3 4 1 7` exists in the `/etc/sysctl.conf` file.

End Point Operations Management Agent

The End Point Operations Management agent adds agent-based discovery and monitoring capabilities to vRealize Operations Manager.

The End Point Operations Management agent is installed on the hosts directly and might or might not be at the same level of trust as the End Point Operations Management server. Therefore, you must verify that the agents are securely installed.

Security Best Practices for Running End Point Operations Management Agents

You must follow certain security best practices while using user accounts.

- For a silent installation, remove any credentials and server certificate thumbprints that were stored in the `AGENT_HOME/conf/agent.properties` file.
- Use a vRealize Operations Manager user account reserved specifically for End Point Operations Management agent registration. For more information, see the topic called "Roles and Privileges" in vRealize Operations Manager in the vRealize Operations Manager Help.
- Disable the vRealize Operations Manager user account that you use for agent registration after the installation is over. You must enable the user's access for agent administration activities. For more information, see the topic called Configuring Users and Groups in vRealize Operations Manager in the vRealize Operations Manager Help.
- If a system that runs an agent is compromised, you can revoke the agent certificate using the vRealize Operations Manager user interface by removing the agent resource. See the section called Revoking an Agent for more detail.

Minimum Required Permissions for Agent Functionality

You require permissions to install and modify a service. If you want to discover a running process, the user account you use to run the agent must also have privileges to access the processes and programs. For Windows operating system installations, you require permissions to install and modify a service. For Linux installations, you require permission to install the agent as a service, if you install the agent using a RPM installer.

The minimum credentials that are required for the agent to register with the vRealize Operations Manager server are those for a user granted the Agent Manager role, without any assignment to objects within the system.

Linux Based Platform Files and Permissions

After you install the End Point Operations Management agent, the owner is the user that installs the agent.

The installation directory and file permissions such as 600 and 700, are set to the owner when the user who installs the End Point Operations Management agent extracts the TAR file or installs the RPM.

Note When you extract the ZIP file, the permissions might not be correctly applied. Verify and ensure that the permissions are correct.

All the files that are created and written to by the agent are given 700 permissions with the owner being the user who runs the agent.

Table 2-16. Linux Files and Permissions

Directory or File	Permissions	Groups or Users	Read	Write	Execute
agent directory/bin	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
agent directory/conf	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
agent directory/log	700	Owner	Yes	Yes	No
		Group	No	No	No
		All	No	No	No
agent directory/data	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
agent directory/bin/ep-agent.bat	600	Owner	Yes	Yes	No
		Group	No	No	No
		All	No	No	No
agent directory/bin/ep-agent.sh	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
agent directory/conf/* (all files in the `conf` directory)	600	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
agent directory/log/* (all files in the `log` directory)	600	Owner	Yes	Yes	No
		Group	No	No	No
		All	No	No	No
agent directory/data/* (all files in the `data` directory)	600	Owner	Yes	Yes	No

Table 2-16. Linux Files and Permissions (continued)

Directory or File	Permissions	Groups or Users	Read	Write	Execute
		Group	No	No	No
		All	No	No	No

Windows Based Platform Files and Permissions

For a Windows based installation of the End Point Operations Management agent, the user installing the agent must have permissions to install and modify the service.

After you install the End Point Operations Management agent, the installation folder including all subdirectories and files should only be accessible by the SYSTEM, the administrators group, and the installation user. When you install the End Point Operations Management agent using `ep-agent.bat`, ensure that the hardening process succeeds. As the user installing the agent, it is advised that you take note of any error messages. If the hardening process fails, the user can apply these permissions manually.

Table 2-17. Windows Files and Permissions

Directory or File	Groups or Users	Full Control	Modify	Read and Execute	Read	Write
<agent directory>/bin	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/conf	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/log	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/data	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/bin/hq-agent.bat	SYSTEM	Yes	-	-	-	-

Table 2-17. Windows Files and Permissions (continued)

Directory or File	Groups or Users	Full Control	Modify	Read and Execute	Read	Write
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/bin/hq-agent.sh	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/conf/* (all files in the `conf` directory)	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/log/* (all files in the `log` directory)	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-
<agent directory>/data/* (all files in `data` directory)	SYSTEM	Yes	-	-	-	-
	Administrator	Yes	-	-	-	-
	Installation User	Yes	-	-	-	-
	Users		-	-	-	-

Open Ports on Agent Host

The agent process listens for commands on two ports 127.0.0.1:2144 and 127.0.0.1:32000 that are configurable. These ports might be arbitrarily assigned, and so, the exact port number might vary. The agent does not open ports on external interfaces.

Table 2-18. Minimum Required Ports

Port	Protocol	Direction	Comments
443	TCP	Outgoing	Used by the agent for outgoing connections over HTTP, TCP, or ICMP.
2144	TCP	Listening	Internal Only. Configurable. Used for inter-process communication between the agent and the command line that loads and configures it. The agent process listens on this port. **Note** The port number is assigned arbitrarily and might differ.
32000	TCP	Listening	Internal Only. Configurable. Used for inter-process communication between the agent and the command line that loads and configures it. The agent process listens on this port. **Note** The port number is assigned arbitrarily and might differ.

Revoking an Agent

If for any reason you need to revoke an agent, for example when a system with a running agent is compromised, you can delete the agent resource from the system. Any subsequent request will fail verification.

Use the vRealize Operations Manager user interface to revoke the agent certificate by removing the agent resource. For more information, see Removing the Agent Resource.

When the system is secured again, you can reinstate the agent. For more information, see Reinstate an Agent Resource.

Removing the Agent Resource

You can use the vRealize Operations Manager to revoke the agent certificate by removing the agent resource.

Prerequisites

To preserve the continuity of the resource with previously recorded metric data, take a record of the End Point Operations Management agent token that is displayed in the resource details.

Procedure

1. Navigate to the Inventory Explorer in the vRealize Operations Manager user interface.
2. Open the Adapter Types tree.
3. Open the EP Ops Adapter list.
4. Select **EP Ops Agent - *HOST_DNS_NAME***.
5. Click **Edit Object**.
6. Record the agent ID, which is the agent token string.
7. Close the Edit Object dialog box .
8. Select **EP Ops Agent - *HOST_DNS_NAME*** and click **Delete Object**.

Reinstate an Agent Resource

When the secure state of a system is recovered, you can reinstate a revoked agent. This ensures that the agent continues to report on the same resources without losing historical data. To do this you must create a new End Point Operations Management token file by using the same token recorded before you removed the agent resource. See the section called Removing The Agent Resource.

Prerequisites

- Ensure that you have the recorded End Point Operations Management token string.
- Use the resource token recorded prior to removing the agent resource from the vRealize Operations Manager server.
- Ensure that you have the Manage Agent privilege.

Procedure

1. Create the agent token file with the user that runs the agent.

   For example, run the command to create a token file containing the 123-456-789 token.

   - On Linux:

     ```
     echo 123-456-789 > /etc/epops-token
     ```

   - On Windows:

     ```
     echo 123-456-789 > %PROGRAMDATA%\VMware\Ep Ops Agent\epops-token
     ```

   In the example, the token file is written to the default token location for that platform

2. Install a new agent and register it with the vRealize Operations Manager server. Ensure that the agent loads the token you inserted in the token file.

   You must have the Manage Agent privilege to perform this action.

Agent Certificate Revocation and Update of Certificates

The reissue flow is initiated from the agent using the `setup` command line argument. When an agent that is already registered uses the `setup` command line argument `ep-agent.sh setup` and fills in the required credentials, a new `registerAgent` command is sent to the server.

The server detects that the agent is already registered and sends the agent a new client certificate without creating another agent resource. On the agent side, the new client certificate replaces the old one. In cases where the server certificate is modified and you run the `ep-agent.sh setup` command, you will see a message that asks you to trust the new certificate. You can alternatively provide the new server certificate thumbprint in the `agent.properties` file prior to running the `ep-agent.sh setup` command, in order to make the process silent.

Prerequisites

Manage agent privilege to revoke and update certificates.

Procedure

- On Linux based operating systems, run the `ep-agent.sh setup` command on the agent host. On Windows based operating systems, run the `ep-agent.bat setup` command.

  If the agent detects that the server certificate has been modified, a message is displayed. Accept the new certificate if you trust it and it is valid.

Patching and Updating the End Point Operations Management Agent

If required, new End Point Operations Management agent bundles are available independent of vRealize Operations Manager releases.

Patches or updates are not provided for the End Point Operations Management agent. You must install the latest available version of the agent that includes the latest security fixes. Critical security fixes will be communicated as per the VMware security advisory guidance. See the topic on Security Advisories.

Additional Secure Configuration Activities

Verify the server user accounts and delete unnecessary applications from the host servers. Block unnecessary ports and disable the services running on your host server that are not required.

Verify Server User Account Settings

It is recommended that you verify that no unnecessary user accounts exist for local and domain user accounts and settings.

Restrict any user account not related to the functioning of the application to those accounts required for administration, maintenance, and troubleshooting. Restrict remote access from domain user accounts to the minimum required to maintain the server. Strictly control and audit these accounts.

Delete and Disable Unnecessary Applications

Delete the unnecessary applications from the host servers. Each additional and unnecessary application increases the risk of exposure because of their unknown or unpatched vulnerabilities.

Disabling Unnecessary Ports and Services

Verify the host server's firewall for the list of open ports that allow traffic.

Block all the ports that are not listed as a minimum requirement for vRealize Operations Manager in the Configuring Ports and Protocols section of this document, or are not required. In addition, audit the services running on your host server and disable those that are not required.

Network Security and Secure Communication

As a security best practice, review and edit the network communication settings of your VMware virtual appliances and host machines. You must also configure the minimum incoming and outgoing ports for vRealize Operations Manager.

Configuring Network Settings for Virtual Application Installation

To ensure that your VMware virtual appliance and host machines allow only safe and essential communication, review and edit their network communication settings.

Prevent User Control of Network Interfaces

As a security best practice, restrict the ability to change the network interface setting to privileged users. If users manipulate network interfaces, it might result in bypassing network security mechanisms or denial of service. Ensure that network interfaces are not configured for user control.

Procedure

1. To verify user control settings, run the `#grep –i '^USERCONTROL=' /etc/sysconfig/network/ifcfg*` command.
2. Make sure that each interface is set to `NO`.

Set the Queue Size for TCP Backlog

As a security best practice, configure a default TCP backlog queue size on VMware appliance host machines. To mitigate TCP denial or service attacks, set an appropriate default size for the TCP backlog queue size. The recommended default setting is 1280.

Procedure

1. Run the `# cat /proc/sys/net/ipv4/tcp_max_syn_backlog` command on each VMware appliance host machine.
2. Set the queue size for TCP backlog.
   a. Open the `/etc/sysctl.conf` file in a text editor.
   b. Set the default TCP backlog queue size by adding the following entry to the file.

      ```
      net.ipv4.tcp_max_syn_backlog=1280
      ```

   c. Save your changes and close the file.

Deny ICMPv4 Echoes to Broadcast Address

Responses to broadcast Internet Control Message Protocol (ICMP) echoes provide an attack vector for amplification attacks and can facilitate network mapping by malicious agents. Configuring your system to ignore ICMPv4 echoes provides protection against such attacks.

Procedure

1. Run the `# cat /proc/sys/net/ipv4/icmp_echo_ignore_broadcasts` command to verify that the system is not sending responses to ICMP broadcast address echo requests.

2 Configure the host system to deny ICMPv4 broadcast address echo requests.

   a Open the `/etc/sysctl.conf` file in a text editor.

   b If the value for this entry is not set to `1`, add the `net.ipv4.icmp_echo_ignore_broadcasts=1` entry.

   c Save the changes and close the file.

Configure the Host System to Disable IPv4 Proxy ARP

IPv4 Proxy ARP allows a system to send responses to ARP requests on one interface on behalf of hosts connected to another interface. You must disable IPv4 Proxy ARP to prevent unauthorized information sharing. Disable the setting to prevent leakage of addressing information between the attached network segments.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/proxy_arp|egrep "default|all"` command to verify whether the Proxy ARP is disabled.

2 Configure the host system to disable IPv4 Proxy ARP.

   a Open the `/etc/sysctl.conf` file in a text editor.

   b If the values are not set to `0`, add the entries or update the existing entries accordingly. Set the value to `0`.

```
net.ipv4.conf.all.proxy_arp=0
net.ipv4.conf.default.proxy_arp=0
```

   c Save any changes you made and close the file.

Configure the Host System to Ignore IPv4 ICMP Redirect Messages

As a security best practice, verify that the host system ignores IPv4 Internet Control Message Protocol (ICMP) redirect messages. A malicious ICMP redirect message can allow a man-in-the-middle attack to occur. Routers use ICMP redirect messages to notify hosts that a more direct route exists for a destination. These messages modify the host's route table and are unauthenticated.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/accept_redirects|egrep "default|all"` command on the host system to check whether the host system ignores IPv4 redirect messages.

2 Configure the host system to ignore IPv4 ICMP redirect messages.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv4.conf.all.accept_redirects=0
net.ipv4.conf.default.accept_redirects=0
```

   c Save the changes and close the file.

Configure the Host System to Ignore IPv6 ICMP Redirect Messages

As a security best practice, verify that the host system ignores IPv6 Internet Control Message Protocol (ICMP) redirect messages. A malicious ICMP redirect message might allow a man-in-the-middle attack to occur. Routers use ICMP redirect messages to tell hosts that a more direct route exists for a destination. These messages modify the host's route table and are unauthenticated.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_redirects|egrep "default|all"` command on the host system and check whether it ignores IPv6 redirect messages.

2 Configure the host system to ignore IPv6 ICMP redirect messages.

   a Open the `/etc/sysctl.conf` to configure the host system to ignore the IPv6 redirect messages.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.accept_redirects=0
net.ipv6.conf.default.accept_redirects=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv4 ICMP Redirects

As a security best practice, verify that the host system denies IPv4 Internet Control Message Protocol (ICMP) redirects. Routers use ICMP redirect messages to inform servers that a direct route exists for a particular destination. These messages contain information from the system's route table that might reveal portions of the network topology.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/send_redirects|egrep "default|all"` on the host system to verify whether it denies IPv4 ICMP redirects.

2 Configure the host system to deny IPv4 ICMP redirects.

   a Open the `/etc/sysctl.conf` file to configure the host system.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv4.conf.all.send_redirects=0
net.ipv4.conf.default.send_redirects=0
```

   c Save the changes and close the file.

Configure the Host System to Log IPv4 Martian Packets

As a security best practice, verify that the host system logs IPv4 Martian packets. Martian packets contain addresses that the system knows to be invalid. Configure the host system to log the messages so that you can identify misconfigurations or attacks in progress.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/log_martians|egrep "default|all"` command to check whether the host logs IPv4 Martian packets.

2 Configure the host system to log IPv4 Martian packets.

   a Open the `/etc/sysctl.conf` file to configure the host system.

   b If the values are not set to `1`, add the following entries to the file or update the existing entries accordingly. Set the value to `1`.

```
net.ipv4.conf.all.log_martians=1
net.ipv4.conf.default.log_martians=1
```

   c Save the changes and close the file.

Configure the Host System to use IPv4 Reverse Path Filtering

As a security best practice, configure your host machines to use IPv4 reverse path filtering. Reverse path filtering protects against spoofed source addresses by causing the system to discard packets with source addresses that have no route or if the route does not point towards the originating interface.

Configure your system to use reverse-path filtering whenever possible. Depending on the system role, reverse-path filtering might cause legitimate traffic to be discarded. In such cases, you might need to use a more permissive mode or disable reverse-path filtering altogether.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/rp_filter|egrep "default|all"` command on the host system to check whether the system uses IPv4 reverse path filtering.

2 Configure the host system to use IPv4 reverse path filtering.

   a Open the `/etc/sysctl.conf` file to configure the host system.

   b If the values are not set to `1`, add the following entries to the file or update the existing entries accordingly. Set the value to `1`.

```
net.ipv4.conf.all.rp_filter=1
net.ipv4.conf.default.rp_filter=1
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv4 Forwarding

As a security best practice, verify that the host system denies IPv4 forwarding. If the system is configured for IP forwarding and is not a designated router, it could be used to bypass network security by providing a path for communication that is not filtered by network devices.

Procedure

1 Run the `# cat /proc/sys/net/ipv4/ip_forward` command to verify whether the host denies IPv4 forwarding.

2 Configure the host system to deny IPv4 forwarding.

   a Open the `/etc/sysctl.conf` to configure the host system.

   b If the value is not set to `0`, add the following entry to the file or update the existing entry accordingly. Set the value to `0`.

```
net.ipv4.ip_forward=0
```

   c Save the changes and close the file.

Configure the Host System to Deny Forwarding of IPv4 Source Routed Packets

Source-routed packets allow the source of the packet to suggest that routers forward the packet along a different path than what is configured on the router, which can be used to bypass network security measures.

This requirement applies only to the forwarding of source-routed traffic, such as when IPv4 forwarding is enabled and the system is functioning as a router.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv4/conf/*/accept_source_route|egrep "default|all"` command to verify whether the system does not use IPv4 source routed packets

2 Configure the host system to deny forwarding of IPv4 source routed packets.

   a Open the `/etc/sysctl.conf` file with a text editor.

   b If the values are not set to `0`, ensure that `net.ipv4.conf.all.accept_source_route=0` and the `et.ipv4.conf.default.accept_source_route=0` are set to `0`.

   c Save and close the file.

Configure the Host System to Deny IPv6 Forwarding

As a security best practice, verify that the host system denies IPv6 forwarding. If the system is configured for IP forwarding and is not a designated router, it can be used to bypass network security by providing a path for communication that is not filtered by network devices.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/forwarding|egrep "default|all"` command to verify whether the host denies IPv6 forwarding.

2 Configure the host system to deny IPv6 forwarding.

   a Open the `/etc/sysctl.conf` to configure the host system.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.forwarding=0
net.ipv6.conf.default.forwarding=0
```

   c Save the changes and close the file.

Configure the Host System to Use IPv4 TCP Syncookies

As a security best practice, verify that the host system uses IPv4 Transmission Control Protocol (TCP) Syncookies. A TCP SYN flood attack might cause a denial of service by filling a system's TCP connection table with connections in the SYN_RCVD state. Syncookies are used so as not to track a connection until a subsequent ACK is received, verifying that the initiator is attempting a valid connection and is not a flood source.

This technique does not operate in a fully standards-compliant manner, but is only activated when a flood condition is detected, and allows defence of the system while continuing to service valid requests.

Procedure

1 Run the `# cat /proc/sys/net/ipv4/tcp_syncookies` command to verify whether the host system uses IPv4 TCP Syncookies.

2 Configure the host system to use IPv4 TCP syncookies.

   a Open the `/etc/sysctl.conf` to configure the host system.

   b If the value is not set to `1`, add the following entry to the file or update the existing entry accordingly. Set the value to `1`.

```
net.ipv4.tcp_syncookies=1
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Advertisements

As a security best practice, verify that the host system denies the acceptance of router advertisements and Internet Control Message Protocol (ICMP) redirects unless necessary. A feature of IPv6 is how systems can configure their networking devices by automatically using information from the network. From a security perspective, it is preferable to manually set important configuration information rather than accepting it from the network in an unauthenticated way.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra|egrep "default|all"` command on the host system to verify whether the system denies the acceptance of router advertisements and ICMP redirects unless necessary.

2 Configure the host system to deny IPv6 router advertisements.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.accept_ra=0
net.ipv6.conf.default.accept_ra=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Solicitations

As a security best practice, verify that host system denies IPv6 router solicitations unless necessary. The router solicitations setting determines how many router solicitations are sent when bringing up the interface. If addresses are assigned statically, there is no need to send any solicitations.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/router_solicitations|egrep "default|all"` command to verify whether the host system denies IPv6 router solicitations unless necessary.

2 Configure the host system to deny IPv6 router solicitations.

   a Open the `/etc/sysctl.conf`.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.router_solicitations=0
net.ipv6.conf.default.router_solicitations=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Preference in Router Solicitations

As a security best practice, verify that your host system denies IPv6 router solicitations unless necessary. The router preference in the solicitations setting determines router preferences. If addresses are assigned statically, there is no need to receive any router preference for solicitations.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_rtr_pref|egrep "default|all"` on the host system to verify whether the host system denies IPv6 router solicitations.

2 Configure the host system to deny IPv6 router preference in router solicitations.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.accept_ra_rtr_pref=0
net.ipv6.conf.default.accept_ra_rtr_pref=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Prefix

As a security best practice, verify that the host system denies IPv6 router prefix information unless necessary. The `accept ra pinfo` setting controls whether the system accepts prefix information from the router. If addresses are statically assigned, the system does not receive any router prefix information.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_pinfo|egrep "default|all"` to verify if that system denies IPv6 router prefix information.

2 Configure the host system to deny IPv6 router prefix.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.accept_ra_pinfo=0
net.ipv6.conf.default.accept_ra_pinfo=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Advertisement Hop Limit Settings

As a security best practice, verify that the host system denies IPv6 router advertisement Hop Limit settings from a router advertisement unless necessary. The `accept_ra_defrtr` setting controls whether the system will accept Hop Limit settings from a router advertisement. Setting it to `0` prevents a router from changing your default IPv6 Hop Limit for outgoing packets.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_defrtr|egrep "default|all"` command to verify that the host system denies IPv6 router Hop Limit settings.

2 If the values are not set to `0`, configure the host system to deny IPv6 router advertisement Hop Limit settings.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.accept_ra_defrtr=0
net.ipv6.conf.default.accept_ra_defrtr=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Router Advertisement Autoconf Settings

As a security best practice, verify that the host system denies IPv6 router advertisement `autoconf` settings. The `autoconf` setting controls whether router advertisements can cause the system to assign a global unicast address to an interface.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/autoconf|egrep "default|all"` command to verify whether the host system denies IPv6 router advertisement `autoconf` settings.

2 If the values are not set to `0`, configure the host system to deny IPv6 router advertisement autoconf settings.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.autoconf=0
net.ipv6.conf.default.autoconf=0
```

   c Save the changes and close the file.

Configure the Host System to Deny IPv6 Neighbor Solicitations

As a security best practice, verify that the host system denies IPv6 neighbor solicitations unless necessary. The `dad_transmits` setting determines how many neighbor solicitations are to be sent out per address including global and link-local, when you bring up an interface to ensure the desired address is unique on the network.

Procedure

1 Run the `# grep [01] /proc/sys/net/ipv6/conf/*/dad_transmits|egrep "default|all"` command to verify whether the host system denies IPv6 neighbor solicitations.

2 If the values are not set to `0`, configure the host system to deny IPv6 neighbor solicitations.

   a Open the `/etc/sysctl.conf` file.

   b If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.dad_transmits=0
net.ipv6.conf.default.dad_transmits=0
```

   c Save the changes and close the file.

Configure the Host System to Restrict IPv6 Maximum Addresses

As a security best practice, verify that the host restricts the maximum number of IPv6 addresses that can be assigned. The maximum addresses setting determines how many global unicast IPv6 addresses can be assigned to each interface. The default is 16 but you must set the number to the statically configured global addresses required.

Procedure

1 Run the `# grep [1] /proc/sys/net/ipv6/conf/*/max_addresses|egrep "default|all"` command to verify whether the host system restricts the maximum number of IPv6 addresses that can be assigned.

2 If the values are not set to 1, configure the host system to restrict the maximum number of IPv6 addresses that can be assigned.

   a Open the `/etc/sysctl.conf` file.

   b Add the following entries to the file or update the existing entries accordingly. Set the value to 1.

```
net.ipv6.conf.all.max_addresses=1
net.ipv6.conf.default.max_addresses=1
```

   c Save the changes and close the file.

Configuring Ports and Protocols

As a security best practice, disable all non-essential ports and protocols.

Configure the minimum incoming and outgoing ports for vRealize Operations Manager components as required for important system components to operate in production.

Minimum Default Incoming Ports

As a security best practice, configure the incoming ports required for vRealize Operations Manager to operate in production.

Table 2-19. Minimum Required Incoming Ports

Port	Protocol	Comments
443	TCP	Used to access the vRealize Operations Manager user interface and the vRealize Operations Manager administrator interface.
123	UDP	Used by vRealize Operations Manager for Network Time Protocol (NTP) synchronization to the master node.
5433	TCP	Used by the master and replica nodes to replicate the global database (vPostgreSQL) when high availability is enabled .
7001	TCP	Used by Cassandra for secure inter-node cluster communication. Do not expose this port to the internet. Add this port to a firewall.
9042	TCP	Used by Cassandra for secure client-related communication among nodes. Do not expose this port to the internet. Add this port to a firewall.
6061	TCP	Used by clients to connect to the GemFire Locator to get connection information to servers in the distributed system. Also monitors server load to send clients to the least-loaded servers.

Table 2-19. Minimum Required Incoming Ports (continued)

Port	Protocol	Comments
10000-10010	TCP and UDP	GemFire Server ephemeral port range used for unicast UDP messaging and for TCP failure detection in a peer-to-peer distributed system.
20000-20010	TCP and UDP	GemFire Locator ephemeral port range used for unicast UDP messaging and for TCP failure detection in a peer-to-peer distributed system.

Table 2-20. Optional Incoming Ports

Port	Protocol	Comments
22	TCP	Optional. Secure Shell (SSH). The SSH service listening on port 22, or any other port, must be disabled in a production environment, and port 22 must be closed.
80	TCP	Optional. Redirects to 443.
3091-3101	TCP	When Horizon View is installed, used to access data for vRealize Operations Manager from Horizon View.

Auditing and Logging on your vRealize Operations Manager System

As a security best practice, set up auditing and logging on your vRealize Operations Manager system.

The detailed implementation of auditing and logging is outside the scope of this document.

Remote logging to a central log host provides a secure store for logs. By collecting log files to a central host, you can easily monitor the environment with a single tool. You can also perform aggregate analysis and search for coordinated attacks on multiple entities within the infrastructure. Logging to a secure, centralized log server can help prevent log tampering and also provide a long-term audit record.

Securing the Remote Logging Server

As a security best practice, ensure that the remote logging server can be configured only by an authorized user and is secure.

Attackers who breach the security of your host machine might search for and attempt to tamper with log files to cover their tracks and maintain control without being discovered.

Use an Authorized NTP Server

Ensure that all the host systems use the same relative time source, including the relevant localization offset. You can correlate the relative time source to an agreed-upon time standard such as Coordinated Universal Time (UTC).

You can easily track and correlate an intruder's actions when you review the relevant log files. Incorrect time settings can make it difficult to inspect and correlate log files to detect attacks, and can make auditing inaccurate. You can use at the least three NTP servers from outside time sources or configure a few local NTP servers on a trusted network that obtain their time from at least three outside time sources.

Client Browser Considerations

As a security best practice, do not use vRealize Operations Manager from untrusted or unpatched clients or from clients that use browser extensions.

Installing

3

You install VMware vRealize Operations Manager to create and configure one or more VMware vRealize Operations Manager nodes that collect and analyze object data rom your enviroment.

This chapter includes the following topics:

- About Installing
- Preparing for Installation
- Installing vRealize Operations Manager
- Resize your Cluster by Adding Nodes
- vRealize Operations Manager Post-Installation Considerations
- Updating, Migrating and Restoring

About Installing

When you install vRealize Operations Manager, you can install the product in an environment that has never been monitored by vRealize Operations Manager. You can also migrate, which captures an environment monitored by a previous version of vRealize Operations Manager so that the new copy of vRealize Operations Manager can monitor that environment.

You can migrate at installation time, or you can postpone a migration until after your copy of vRealize Operations Manager is in production use. In other words, you can run vRealize Operations Manager to monitor a fresh environment, and at any time, decide to add an environment that was being monitored by a previous vRealize Operations Manager.

Installation Overview

You prepare for vRealize Operations Manager installation by evaluating your environment and deploying enough vRealize Operations Manager cluster nodes to support how you want to use the product.

Workflow of vRealize Operations Manager Installation

The vRealize Operations Manager virtual appliance installation process consists of deploying the vRealize Operations Manager OVF or an installer once for each cluster node, accessing the

product to set up cluster nodes according to their role, and logging in to configure the installation.

Figure 3-1. vRealize Operations Manager Installation Architecture

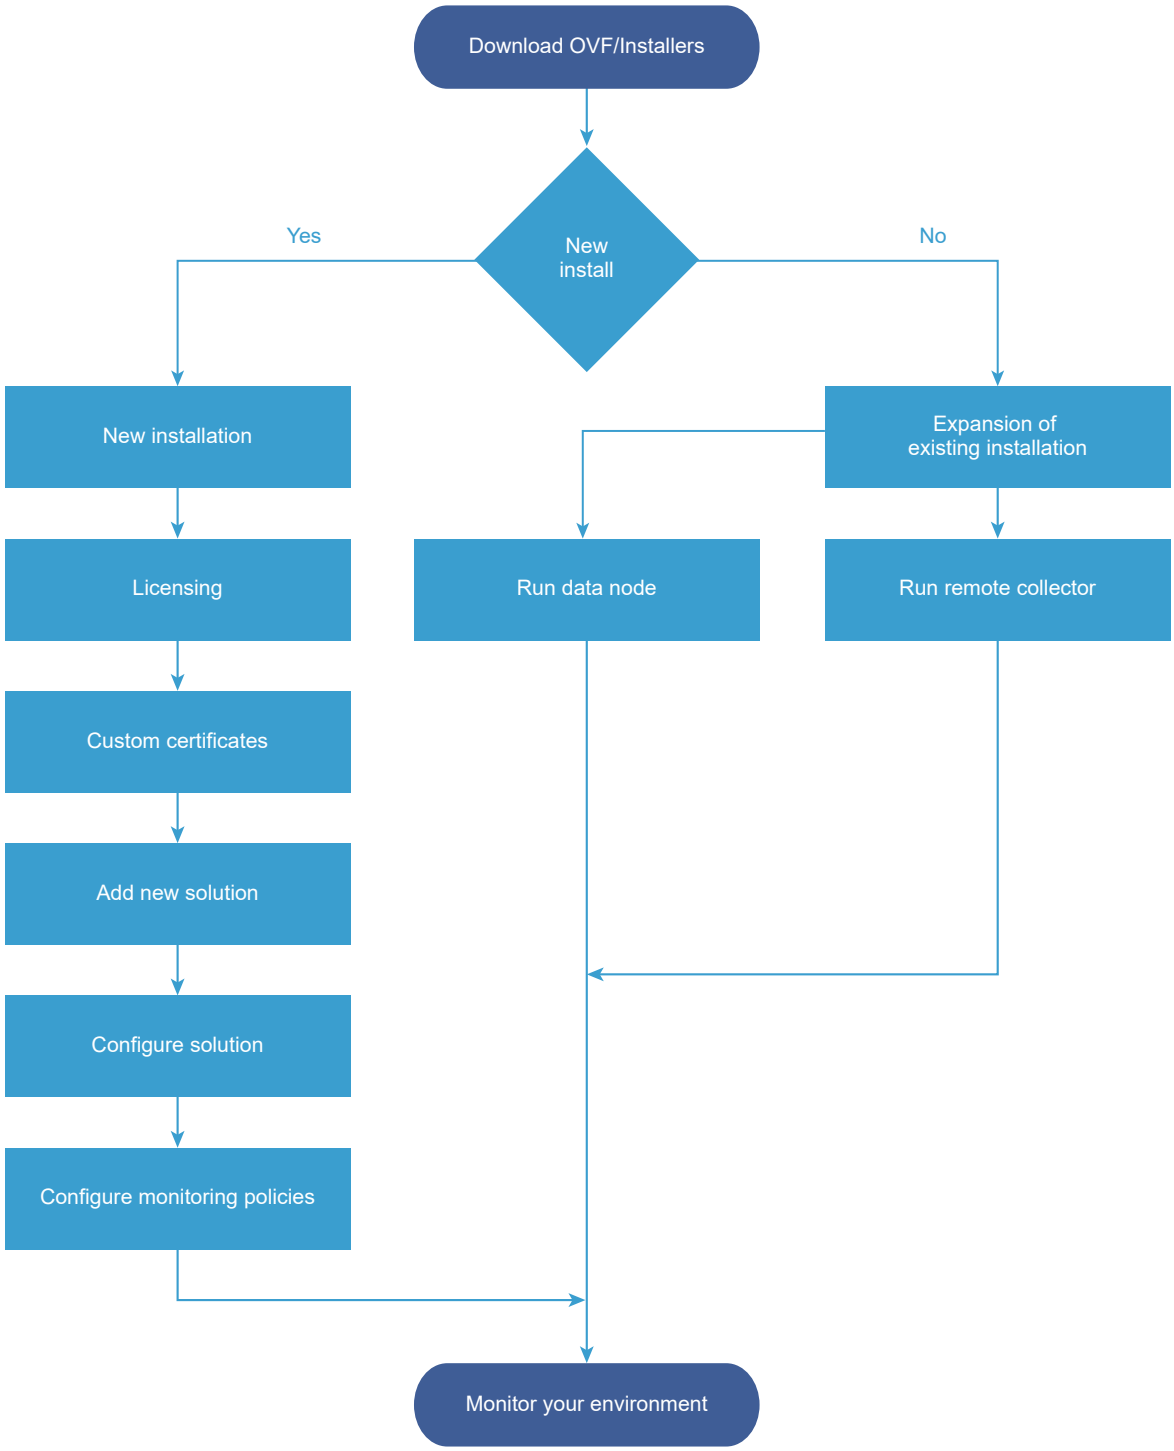

Sizing the vRealize Operations Manager Cluster

The resources needed for vRealize Operations Manager depend on how large of an environment you expect to monitor and analyze, how many metrics you plan to collect, and how long you need to store the data.

It is difficult to broadly predict the CPU, memory, and disk requirements that will meet the needs of a particular environment. There are many variables, such as the number and type of objects collected, which includes the number and type of adapters installed, the presence of HA, the duration of data retention, and the quantity of specific data points of interest, such as symptoms, changes, and so on.

VMware expects vRealize Operations Manager 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.

Knowledge Base article 2093783

The Knowledge Base articles include overall maximums, plus spreadsheet calculators in which you enter the number of objects and metrics that you expect to monitor. To obtain the numbers, some users take the following high-level approach, which uses vRealize Operations Manager itself.

1 Review this guide to understand how to deploy and configure a vRealize Operations Manager node.

2 Deploy a temporary vRealize Operations Manager node.

3 Configure one or more adapters, and allow the temporary node to collect overnight.

4 Access the Cluster Management page on the temporary node.

5 Using the Adapter Instances list in the lower portion of the display as a reference, enter object and metric totals of the different adapter types into the appropriate sizing spreadsheet from Knowledge Base article 2093783.

6 Deploy the vRealize Operations Manager cluster based on the spreadsheet sizing recommendation. You can build the cluster by adding resources and data nodes to the temporary node or by starting over.

If you have a large number of adapters, you might need to reset and repeat the process on the temporary node until you have all the totals you need. The temporary node will not have enough capacity to simultaneously run every connection from a large enterprise.

Another approach to sizing is through self monitoring. Deploy the cluster based on your best estimate, but create an alert for when capacity falls below a threshold, one that allows enough time to add nodes or disk to the cluster. You also have the option to create an email notification when thresholds are passed.

During internal testing, a single-node vApp deployment of vRealize Operations Manager that monitored 8,000 virtual machines ran out of disk storage within one week.

Add Data Disk Space to a vRealize Operations Manager vApp Node

You add to the data disk of vRealize Operations Manager vApp nodes when space for storing the collected data runs low.

Prerequisites

- Note the disk size of the analytics cluster nodes. When adding disk, you must maintain uniform size across analytics cluster nodes.
- Use the vRealize Operations Manager administration interface to take the node offline.
- Verify that you are connected to a vCenter Server system with a vSphere client, and log in to the vSphere client.

Procedure

1. Shut down the virtual machine for the node.
2. Edit the hardware settings of the virtual machine, and add another disk.

   Note Do not expand disks. vRealize Operations Manager does not support expanding disks.

3. Power on the virtual machine for the node.

Results

During the power-on process, the virtual machine expands the vRealize Operations Manager data partition.

Complexity of Your Environment

When you deploy vRealize Operations Manager, the number and nature of the objects that you want to monitor might be complex enough to recommend a Professional Services engagement.

Complexity Levels

Every enterprise is different in terms of the systems that are present and the level of experience of deployment personnel. The following table presents a color-coded guide to help you determine where you are on the complexity scale.

- Green

  Your installation only includes conditions that most users can understand and work with, without assistance. Continue your deployment.

- Yellow

  Your installation includes conditions that might justify help with your deployment, depending on your level of experience. Consult your account representative before proceeding, and discuss using Professional Services.

- Red

Your installation includes conditions that strongly recommend a Professional Services engagement. Consult your account representative before proceeding, and discuss using Professional Services.

Note that these color-coded levels are not firm rules. Your product experience, which increases as you work with vRealize Operations Manager and in partnership with Professional Services, must be taken into account when deploying vRealize Operations Manager.

Table 3-1. Effect of Deployment Conditions on Complexity

Complexity Level	Current or New Deployment Condition	Additional Notes
Green	You run only one vRealize Operations Manager deployment.	Lone instances are usually easy to create in vRealize Operations Manager.
Green	Your deployment includes a management pack that is listed as Green according to the compatibility guide on the VMware Solutions Exchange Web site.	The compatibility guide indicates whether the supported management pack for vRealize Operations Manager is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected. Note that the terms *solution*, *management pack*, *adapter*, and *plug-in* are used somewhat interchangeably.
Yellow	You run multiple instances of vRealize Operations Manager.	Multiple instances are typically used to address scaling or operator use patterns.
Yellow	Your deployment includes a management pack that is listed as Yellow according to the compatibility guide on the VMware Solutions Exchange Web site.	The compatibility guide indicates whether the supported management pack for vRealize Operations Manager is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.
Yellow	You are deploying vRealize Operations Manager remote collector nodes.	Remote collector nodes gather data but leave the storage and processing of the data to the analytics cluster.
Yellow	You are deploying a multiple-node vRealize Operations Manager cluster.	Multiple nodes are typically used for scaling out the monitoring capability of vRealize Operations Manager.
Yellow	Your new vRealize Operations Manager instance will include a Linux based deployment.	Linux deployments are not as common as vApp deployments and often need special consideration.
Yellow	Your vRealize Operations Manager instance will use high availability (HA).	High availability and its node failover capability is a unique multiple-node feature that you might want additional help in understanding.

Table 3-1. Effect of Deployment Conditions on Complexity (continued)

Complexity Level	Current or New Deployment Condition	Additional Notes
Yellow	You want help in understanding the new or changed features in vRealize Operations Manager and how to use them in your environment.	vRealize Operations Manager is different than vCenter Operations Manager in areas such as policies, alerts, compliance, custom reporting, or badges. In addition, vRealize Operations Manager uses one consolidated interface.
Red	You run multiple instances of vRealize Operations Manager, where at least one includes virtual desktop infrastructure (VDI).	Multiple instances are typically used to address scaling, operator use patterns, or because separate VDI (V4V monitoring) and non-VDI instances are needed.
Red	Your deployment includes a management pack that is listed as Red according to the compatibility guide on the VMware Solutions Exchange Web site.	The compatibility guide indicates whether the supported management pack for vRealize Operations Manager is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.
Red	You are deploying multiple vRealize Operations Manager clusters.	Multiple clusters are typically used to isolate business operations or functions.
Red	Your current vRealize Operations Manager deployment required a Professional Services engagement to install it.	If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version.
Red	Professional Services customized your vRealize Operations Manager deployment. Examples of customization include special integrations, scripting, nonstandard configurations, multiple level alerting, or custom reporting.	If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version.

About vRealize Operations Manager Cluster Nodes

All vRealize Operations Manager clusters consist of a master node, an optional replica node for high availability, optional data nodes, and optional remote collector nodes.

When you install vRealize Operations Manager, you use a vRealize Operations Manager vApp deployment to create role-less nodes. After the nodes are created and have their names and IP addresses, you use an administration interface to configure them according to their role.

You can create role-less nodes all at once or as needed. A common as-needed practice might be to add nodes to scale out vRealize Operations Manager to monitor an environment as the environment grows larger.

The following node types make up the vRealize Operations Manager analytics cluster:

Master Node

The initial, required node in vRealize Operations Manager. All other nodes are managed by the master node.

In a single-node installation, the master node manages itself, has adapters installed on it, and performs all data collection and analysis.

Data Node

In larger deployments, additional data nodes have adapters installed and perform collection and analysis.

Larger deployments usually include adapters only on the data nodes so that master and replica node resources can be dedicated to cluster management.

Replica Node

To use vRealize Operations Manager high availability (HA), the cluster requires that you convert a data node into a replica of the master node.

The following node type is a member of the vRealize Operations Manager cluster but not part of the analytics cluster:

Remote Collector Node

Distributed deployments might require a remote collector node that can navigate firewalls, interface with a remote data source, reduce bandwidth across data centers, or reduce the load on the vRealize Operations Manager analytics cluster. Remote collectors only gather objects for the inventory, without storing data or performing analysis. In addition, remote collector nodes may be installed on a different operating system than the rest of the cluster.

About vRealize Operations Manager Remote Collector Nodes

A remote collector node is an additional cluster node that allows vRealize Operations Manager to gather more objects into its inventory for monitoring. Unlike data nodes, remote collector nodes only include the collector role of vRealize Operations Manager, without storing data or processing any analytics functions.

A remote collector node is usually deployed to navigate firewalls, reduce bandwidth across data centers, connect to remote data sources, or reduce the load on the vRealize Operations Manager analytics cluster.

Remote collectors do not buffer data while the network is experiencing a problem. If the connection between remote collector and analytics cluster is lost, the remote collector does not store data points that occur during that time. In turn, and after the connection is restored, vRealize Operations Manager does not retroactively incorporate associated events from that time into any monitoring or analysis.

You must have at least a master node before adding remote collector nodes.

About vRealize Operations Manager High Availability

vRealize Operations Manager supports high availability (HA). HA creates a replica for the vRealize Operations Manager master node and protects the analytics cluster against the loss of a node.

With HA, data stored on the master node is always 100% backed up on the replica node. To enable HA, you must have at least one data node deployed, in addition to the master node.

- HA is not a disaster recovery mechanism. HA protects the analytics cluster against the loss of only one node, and because only one loss is supported, you cannot stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When HA is enabled, the replica can take over all functions that the master provides, were the master to fail for any reason. If the master fails, failover to the replica is automatic and requires only two to three minutes of vRealize Operations Manager downtime to resume operations and restart data collection.

  When a master node problem causes failover, the replica node becomes the master node, and the cluster runs in degraded mode. To get out of degraded mode, take one of the following steps.

  - Return to HA mode by correcting the problem with the master node. When a master node exits an HA-enabled cluster, master node does not rejoin with the cluster without manual intervention. Therefore, restart the vRealize Operations Analytics process on the downed node to change its role to replica and rejoin the cluster.
  - Return to HA mode by converting a data node into a new replica node and then removing the old, failed master node. Removed master nodes cannot be repaired and re-added to vRealize Operations Manager.
  - Change to non-HA operation by disabling HA and then removing the old, failed master node. Removed master nodes cannot be repaired and re-added to vRealize Operations Manager.
- In the administration interface, after an HA replica node takes over and becomes the new master node, you cannot remove the previous, offline master node from the cluster. In addition, the previous node continues to be listed as a master node. To refresh the display and enable removal of the node, refresh the browser.
- When HA is enabled, the cluster can survive the loss of one data node without losing any data. However, HA protects against the loss of only one node at a time, of any kind, so simultaneously losing data and master/replica nodes, or two or more data nodes, is not supported. Instead, vRealize Operations Manager HA provides additional application level data protection to ensure application level availability.
- When HA is enabled, it lowers vRealize Operations Manager capacity and processing by half, because HA creates a redundant copy of data throughout the cluster, as well as the replica backup of the master node. Consider your potential use of HA when planning the number and size of your vRealize Operations Manager cluster nodes. See Sizing the vRealize Operations Manager Cluster.

- When HA is enabled, deploy analytics cluster nodes on separate hosts for redundancy and isolation. One option is to use anti-affinity rules that keep nodes on specific hosts in the vSphere cluster.

  If you cannot keep the nodes separate, you should not enable HA. A host fault would cause the loss of more than one node, which is not supported, and all of vRealize Operations Manager would become unavailable.

  The opposite is also true. Without HA, you could keep nodes on the same host, and it would not make a difference. Without HA, the loss of even one node would make all of vRealize Operations Manager unavailable.
- When you power off the data node and change the network settings of the VM, this affects the IP address of the data node. After this point, the HA cluster is no longer accessible and all the nodes have a status of "Waiting for analytics". Verify that you have used a static IP address.
- When you remove a node that has one or more vCenter adapters configured to collect data from a HA-enabled cluster, one or more vCenter adapters associated with that node stops collecting. You change the adapter configuration to pin them to another node before removing the node.
- Administration UI shows the resource cache count, which is created for active objects only, but the Inventory Explorer displays all objects. Therefore, when you remove a node from a HA-enabled cluster allowing the vCenter adapters collect data and rebalance each node, the Inventory explorer displays a different quantity of objects from that shown in the Administration UI.

Preparing for Installation

You preparing for your installation, consider some of these best practises, platform, and cluster requirements.

Requirements

You have to consider important requirements while creating nodes in a vRealize Operations Manager.

Using IPv6 with vRealize Operations Manager

vRealize Operations Manager supports Internet Protocol version 6 (IPv6), the network addressing convention that will eventually replace IPv4. Use of IPv6 with vRealize Operations Manager requires that certain limitations be observed.

Using IPv6

- All vRealize Operations Manager cluster nodes, including remote collectors, must have IPv6 addresses. Do not mix IPv6 and IPv4.

- All vRealize Operations Manager cluster nodes, including remote collectors, must be vApp based.
- Use global IPv6 addresses only. Link-local addresses are not supported.
- If any nodes use DHCP, your DHCP server must be configured to support IPv6.
- DHCP is only supported on data nodes and remote collectors. Master nodes and replica nodes still require fixed addresses, which is true for IPv4 as well.
- Your DNS server must be configured to support IPv6.
- When adding nodes to the cluster, remember to enter the IPv6 address of the master node.
- When registering a VMware vCenter instance within vRealize Operations Manager, place square brackets around the IPv6 address of your VMware vCenter Server system if vCenter is also using IPv6.

  For example: `[2015:0db8:85a3:0042:1000:8a2e:0360:7334]`

  Note that, even when vRealize Operations Manager is using IPv6, vCenter Server may still have an IPv4 address. In that case, vRealize Operations Manager does not need the square brackets.

Cluster Requirements

When you create the cluster nodes that make up vRealize Operations Manager, you have general requirements that you must meet.

General vRealize Operations Manager Cluster Node Requirements

You have to follow some general requirements to create a node on your environment.

General Requirements

- vRealize Operations Manager version. All nodes must run the same vRealize Operations Manager version.

  For example, do not add a version 6.1 data node to a cluster of vRealize Operations Manager 6.2 nodes.
- Analytics Cluster Deployment Type. In the analytics cluster, all nodes must be the same kind of deployment: vApp.
- Remote Collector Deployment Type. A remote collector node does not need to be the same deployment type as the analytics cluster nodes.

  When you add a remote collector of a different deployment type, the following clusters are supported:

  - vApp analytics cluster
- Analytics Cluster Node Sizing. In the analytics cluster, CPU, memory, and disk size must be identical for all nodes.

  Master, replica, and data nodes must be uniform in sizing.

- Remote Collector Node Sizing. Remote collector nodes may be of different sizes from each other or from the uniform analytics cluster node size.
- Geographical Proximity. You may place analytics cluster nodes in different vSphere clusters, but the nodes must reside in the same geographical location.

  Different geographical locations are not supported.
- Virtual Machine Maintenance. When any node is a virtual machine, you may only update the virtual machine software by directly updating the vRealize Operations Manager software.

  For example, going outside of vRealize Operations Manager to access vSphere to update VMware Tools is not supported.
- Redundancy and Isolation. If you expect to enable HA, place analytics cluster nodes on separate hosts. See About vRealize Operations Manager High Availability .
- You can deploy remote collectors behind a firewall. You cannot use NAT between remote collectors and analytics nodes.

Requirements for Solutions

Be aware that solutions might have requirements beyond those for vRealize Operations Manager itself. For example, vRealize Operations Manager for Horizon View has specific sizing guidelines for its remote collectors.

See your solution documentation, and verify any additional requirements before installing solutions. Note that the terms *solution*, *management pack*, *adapter*, and *plug-in* are used somewhat interchangeably.

vRealize Operations Manager Cluster Node Networking Requirements

When you create the cluster nodes that make up vRealize Operations Manager, the associated setup within your network environment is critical to inter-node communication and proper operation.

Networking Requirements

Important vRealize Operations Manager analytics cluster nodes need frequent communication with one another. In general, your underlying vSphere architecture might create conditions where some vSphere actions affect that communication. Examples include, but are not limited to, vMotions, storage vMotions, HA events, and DRS events.

- The master and replica nodes must be use static IP address, or fully qualified domain name (FQDN) with a static IP address.

  Data and remote collector nodes can use dynamic host control protocol (DHCP).
- You can successfully reverse-DNS all nodes, including remote collectors, to their FQDN, currently the node hostname.

  Nodes deployed by OVF have their hostnames set to the retrieved FQDN by default.
- All nodes, including remote collectors, must be bidirectionally routable by IP address or FQDN.

- Do not separate analytics cluster nodes with network address translation (NAT), load balancer, firewall, or a proxy that inhibits bidirectional communication by IP address or FQDN
- Analytics cluster nodes must not have the same hostname.
- Place analytics cluster nodes within the same data center and connect them to the same local area network (LAN).
- Place analytics cluster nodes on same Layer 2 network and IP subnet.

  A stretched Layer 2 or routed Layer 3 network is not supported.
- Do not span the Layer 2 network across sites, which might create network partitions or network performance issues.
- One-way latency between the analytics cluster nodes must be 5 ms or lower.
- Network bandwidth between the analytics cluster nodes must be one gbps or higher.
- Do not distribute analytics cluster nodes over a wide area network (WAN).

  To collect data from a WAN, a remote or separate data center, or a different geographic location, use remote collectors.
- Remote collectors are supported through a routed network but not through NAT.
- Do not include an underscore in the hostname of any cluster node.

vRealize Operations Manager Cluster Node Best Practices

When you create the cluster nodes that make up vRealize Operations Manager, additional best practices improve performance and reliability in vRealize Operations Manager.

Best Practices

- Deploy vRealize Operations Manager analytics cluster nodes in the same vSphere cluster in a single datacenter and add only one node at a time to a cluster allowing it to complete before adding another node.
- If you deploy analytics cluster nodes in a highly consolidated vSphere cluster, you might need resource reservations for optimal performance.

  Determine whether the virtual to physical CPU ratio is affecting performance by reviewing CPU ready time and co-stop.
- Deploy analytics cluster nodes on the same type of storage tier.
- To continue to meet analytics cluster node size and performance requirements, apply storage DRS anti-affinity rules so that nodes are on separate datastores.
- To prevent unintentional migration of nodes, set storage DRS to manual.
- To ensure balanced performance from analytics cluster nodes, use ESXi hosts with the same processor frequencies. Mixed frequencies and physical core counts might affect analytics cluster performance.

- To avoid a performance decrease, vRealize Operations Manager analytics cluster nodes need guaranteed resources when running at scale. The vRealize Operations Manager Knowledge Base includes sizing spreadsheets that calculate resources based on the number of objects and metrics that you expect to monitor, use of HA, and so on. When sizing, it is better to over-allocate than under-allocate resources.

  See Knowledge Base article 2093783.

- Because nodes might change roles, avoid machine names such as Master, Data, Replica, and so on. Examples of changed roles might include making a data node into a replica for HA, or having a replica take over the master node role.

- The NUMA placement is removed in the vRealize Operations Manager 6.3 and later. Procedures related to NUMA settings from the OVA file follow:

  Table 3-2. NUMA Setting

Action	Description
Set the vRealize Operations Manager cluster status to offline	1 Shut down the vRealize Operations Manager cluster. 2 Right-click the cluster and click **Edit Settings > Options > Advanced General**. 3 Click **Configuration Parameters**. In the vSphere Client, repeat these steps for each VM.
Remove the NUMA setting	1 From the Configuration Parameters, remove the setting `numa.vcpu.preferHT` and click **OK**. 2 Click **OK**. 3 Repeat these steps for all the VMs in the vRealize Operations cluster. 4 Power on the cluster.

  Note To ensure the availability of adequate resources and continued product performance, monitor vRealize Operations performance by checking its CPU usage, CPU ready and CPU contention time.

Sizing and Scaling Requirements

The CPU, memory, and disk requirements that meet the needs of a particular environment depend on the number and type of objects in your environment and the data collected. This includes the number and type of adapters installed, the use of HA (High Availability), the duration of data retention, and the quantity of specific data points of interest.

VMware updates Knowledge Base article 2093783 with the most current information about sizing and scaling. The Knowledge Base article includes overall maximums and spreadsheet calculations that provide a recommendation based on the number of objects and metrics you expect to monitor.

Installing vRealize Operations Manager

vRealize Operations Manager nodes are virtual appliance (vApp) based systems.

Deployment of vRealize Operations Manager

vRealize Operations Manager consists of one or more nodes in a cluster. To create these nodes, you have to download and install the vRealize Operations Manager suitable to your environment.

OVF file

vRealize Operations Manager consists of one or more nodes, in a cluster. To create nodes, you use the vSphere client to download and deploy the vRealize Operations Manager virtual machine, once for each cluster node.

Create a Node by Deploying an OVF

vRealize Operations Manager consists of one or more nodes, in a cluster. To create nodes, you use the vSphere client to download and deploy the vRealize Operations Manager virtual machine, once for each cluster node.

Prerequisites

- Verify that you have permissions to deploy OVF templates to the inventory.
- If the ESXi host is part of a cluster, enable DRS in the cluster. If an ESXi host belongs to a non-DRS cluster, all resource pool functions are disabled.
- If this node is to be the master node, reserve a static IP address for the virtual machine, and know the associated domain name server, default gateway, and network mask values.

  Plan to keep the IP address because it is difficult to change the address after installation.
- If this node is to be a data node that will become the HA replica node, reserve a static IP address for the virtual machine, and know the associated domain name server, default gateway, and network mask values.

  In addition, familiarize yourself with HA node placement as described in About vRealize Operations Manager High Availability .
- Preplan your domain and machine naming so that the deployed virtual machine name will begin and end with alphabet (a–z) or digit (0–9) characters, and will only contain alphabet, digit, or hyphen (-) characters. The underscore character (_) must not appear in the host name or anywhere in the fully qualified domain name (FQDN).

  Plan to keep the name because it is difficult to change the name after installation.

  For more information, review the host name specifications from the Internet Engineering Task Force. See www.ietf.org.
- Preplan node placement and networking to meet the requirements described in General vRealize Operations Manager Cluster Node Requirements and vRealize Operations Manager Cluster Node Networking Requirements.
- If you expect the vRealize Operations Manager cluster to use IPv6 addresses, review the IPv6 limitations described in Using IPv6 with vRealize Operations Manager.

- Download the vRealize Operations Manager `.ova` file to a location that is accessible to the vSphere client.
- If you download the virtual machine and the file extension is `.tar`, change the file extension to `.ova`.
- Verify that you are connected to a vCenter Server system with a vSphere client, and log in to the vSphere client.

  Do not deploy vRealize Operations Manager from an ESXi host. Deploy only from vCenter Server.

Procedure

1. Select the vSphere **Deploy OVF Template** option.
2. Enter the path to the vRealize Operations Manager `.ova` file.
3. Follow the prompts until you are asked to enter a name for the node.
4. Enter a node name. Examples might include `Ops1`, `Ops2` or `Ops-A`, `Ops-B`.

   Do not include nonstandard characters such as underscores (_) in node names.

   Use a different name for each vRealize Operations Manager node.
5. Follow the prompts until you are asked to select a configuration size.
6. Select the size configuration that you need. Your selection does not affect disk size.

   Default disk space is allocated regardless of which size you select. If you need additional space to accommodate the expected data, add more disk after deploying the vApp.
7. Follow the prompts until you are asked to select the disk format.

Option	Description
Thick Provision Lazy Zeroed	Creates a virtual disk in a default thick format.
Thick Provision Eager Zeroed	Creates a type of thick virtual disk that supports clustering features such as Fault Tolerance. Thick provisioned eager-zeroed format can improve performance depending on the underlying storage subsystem. Select the thick provisioned eager-zero option when possible.
Thin Provision	Creates a disk in thin format. Use this format to save storage space.

   Snapshots can negatively affect the performance of a virtual machine and typically result in a 25–30 percent degradation for the vRealize Operations Manager workload. Do not use snapshots.
8. Click **Next**.
9. From the drop-down menu, select a Destination Network, for example, **Network 1 = TEST**, and click **Next**.

10 In Properties, under Application, Timezone Setting, leave the default of UTC or select a time zone.

   The preferred approach is to standardize on UTC. Alternatively, configure all nodes to the same time zone.

11 (Optional) Select the option for IPv6.

12 Under Networking Properties, leave the entries blank for DHCP, or fill in the default gateway, domain name server, static IP address, and network mask values.

   The master node and replica node require a static IP. A data node or remote collector node may use DHCP or static IP.

13 Click **Next**.

14 Review the settings and click **Finish**.

15 If you are creating a multiple-node vRealize Operations Manager cluster, repeat through all the steps to deploy each node.

What to do next

Use a Web browser client to configure a newly added node as the vRealize Operations Manager master node, a data node, a high availability master replica node, or a remote collector node. The master node is required first.

Caution For security, do not access vRealize Operations Manager from untrusted or unpatched clients, or from clients using browser extensions.

Installation Types

After you have installed vRealize Operations Manager product, you can either perform a new installation, an express installation or expand an existing installation.

- Express Installation
- New installation
- Expand Installation

Figure 3-2. Getting Started Setup

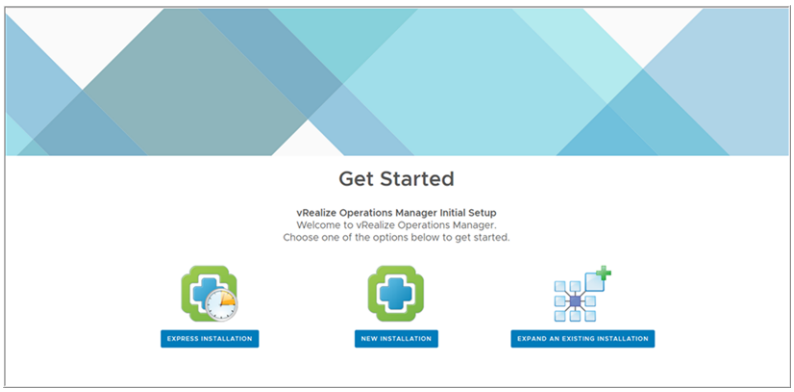

Installing vRealize Operations Manager for a New User

After you install vRealize Operations Manager using an OVF or an installer, you are notified to the main product UI page. You can create a single node or multiple nodes depending on your environment.

Introduction to a New Installation

You can perform a new installation as a first time user and create a single node to handle both administration and data handling.

Figure 3-3. New Installation from the Setup screen

Perform a New Installation on the vRealize Operations Manager product UI

You can create a single node and configure this as a master node or create a master node in a cluster to handle additional data. All vRealize Operations Manager installations require a master node. With a single node cluster, administration and data functions are on the same master node. A multiple-node vRealize Operations Manager cluster contains one master node and one or more nodes for handling additional data.

Prerequisites

- Create a node by deploying the vRealize Operations Manager vApp.
- After it is deployed, note the fully qualified domain name (FQDN) or IP address of the node.
- If you plan to use a custom authentication certificate, verify that your certificate file meets the requirements for vRealize Operations Manager.

Procedure

1. Navigate to the name or IP address of the node that will be the master node of vRealize Operations Manager.

   The setup wizard appears, and you do not need to log in to vRealize Operations Manager.
2. Click **New Installation**.
3. Click **Next**.
4. Enter and confirm a password for the admin user account, and click **Next**.

   Passwords require a minimum of 8 characters, one uppercase letter, one lowercase letter, one digit, and one special character.

   The user account name is admin by default and cannot be changed.

5 Select whether to use the certificate included with vRealize Operations Manager or to install one of your own.

   a To use your own certificate, click **Browse**, locate the certificate file, and click **Open** to load the file in the Certificate Information text box.

   b Review the information detected from your certificate to verify that it meets the requirements for vRealize Operations Manager.

6 Click **Next**.

7 Enter a name for the master node.

   For example: `Ops-Master`

8 Enter the URL or IP address for the Network Time Protocol (NTP) server with which the cluster will synchronize.

   For example: `nist.time.gov`

9 Click **Add**.

   Leave the NTP blank to have vRealize Operations Manager manage its own synchronization by having all nodes synchronize with the master node and replica node.

10 Click **Next**, and click **Finish**.

   The administration interface appears, and it takes a moment for vRealize Operations Manager to finish adding the master node.

Results

You have created a master node to which you can add more nodes.

What to do next

After creating the master node, you have the following options.

- Create and add data nodes to the unstarted cluster.
- Create and add remote collector nodes to the unstarted cluster.
- Click **Start vRealize Operations Manager** to start the single-node cluster, and log in to finish configuring the product.

  The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

About the vRealize Operations Manager Master Node

The master node is the required, initial node in your vRealize Operations Manager cluster.

The master node performs administration for the cluster and must be online before you configure any new nodes. In addition, the master node must be online before other nodes are brought online. If the master node and replica node go offline together, bring them back online separately. Bring the master node completely online first, and then bring the replica node online. For example, if the entire cluster were offline for any reason, you would bring the master node online first.

Advantages of a New installation

You can use the new installation to create a new master node during the first installation of vRealize Operations Manager. With the master node in place, you can then start adding more nodes to form a cluster and then define an environment for your organization.

In a single-node clusters, administration and data are on the same master node. A multiple-node cluster includes one master node and one or more data nodes. In addition, there might be remote collector nodes, and there might be one replica node used for high availability. For more information on creating a master node, see About the vRealize Operations Manager Master Node.

Installing vRealize Operations Manager as an Administrator

As an administrator, you can install several instances of vRealize Operations Manager build in your VM environment.

Introduction to Express Installation

Express installation is one possible way to create master nodes, add data nodes, form clusters and test your connection status. You can use express installation to save time and speed up the process of installation when compared to new installation. It is recommended not to use this feature unless the user is an administrator.

Figure 3-4. Express Installation from the Setup screen

Perform an Express Installation on the vRealize Operations Manager product UI

Use express installation on the vRealize Operations Manager cluster to create a master node. Select express installation option when installing for the first time.

Prerequisites

Verify that you have a static IP address created from either an OVF file.

Procedure

1. Navigate to the name or IP address of the node that will be the master node of vRealize Operations Manager.

   The setup wizard appears, and you do not need to log in to vRealize Operations Manager.
2. Click **Express Installation**.
3. Click **Next**.
4. Enter and confirm a password for the admin user account, and click **Next**.

   Passwords require a minimum of 8 characters, one uppercase letter, one lowercase letter, one digit, and one special character.

   The user account name is admin by default and cannot be changed.
5. Click **Next**.
6. Click **Finish**.

Results

You have created a master node to which you can add more nodes.

Advantages of an Express Installation

Express installation saves time when compared to a new installation to create a new master node. The express installation uses the default certificates, which differs from one organization to another. This feature is mainly used by the developers or the administrators.

Expand an Existing Installation of vRealize Operations Manager

Use this option to add a node to an existing vRealize Operations Manager cluster. You can use this option if you have already configured a master node and you want to increase the capacity by adding more nodes to your cluster.

Introduction to Expand an Existing Installation

You can deploy and configure additional nodes so that vRealize Operations Manager can support larger environments. A master node always requires an additional node for a cluster to monitor your environment. With expanding your installation, you can add more than one node to your cluster.

Adding Data Nodes

Data nodes are the additional cluster nodes that allow you to scale out vRealize Operations Manager to monitor larger environments.

You can dynamically scale out vRealize Operations Manager by adding data nodes without stopping the vRealize Operations Manager cluster. When you scale out the cluster by 25% or more, you should restart the cluster to allow vRealize Operations Manager to update its storage size, and you might notice a decrease in performance until you restart. A maintenance interval provides a good opportunity to restart the vRealize Operations Manager cluster.

In addition, the product administration options include an option to re-balance the cluster, which can be done without restarting. Rebalancing adjusts the vRealize Operations Manager workload across the cluster nodes.

Figure 3-5. Expand an existing installation from the Setup screen

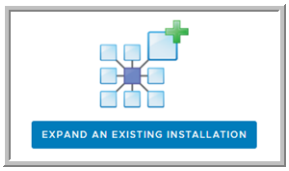

Note Do not shut down online cluster nodes externally or by using any means other than the vRealize Operations Manager interface. Shut down a node externally only after taking it offline in the vRealize Operations Manager interface.

Expand an Existing Installation to Add a Data Node

Larger environments with multiple-node vRealize Operations Manager clusters contain one master node and one or more data nodes for additional data collection, storage, processing, and analysis.

Prerequisites

- Create nodes by deploying the vRealize Operations Manager vApp.
- Create and configure the master node.
- Note the fully qualified domain name (FQDN) or IP address of the master node.

Procedure

1. In a Web browser, navigate to the name or IP address of the node that will become the data node.

   The setup wizard appears, and you do not need to log in to vRealize Operations Manager.
2. Click **Expand an Existing Installation**.
3. Click **Next**.
4. Enter a name for the node (for example, **Data–1**).
5. From the Node Type drop-down, select **Data**.
6. Enter the FQDN or IP address of the master node and click **Validate**.
7. Select **Accept this certificate** and click **Next**.

   If necessary, locate the certificate on the master node and verify the thumbprint.
8. Verify the vRealize Operations Manager administrator username of admin.
9. Enter the vRealize Operations Manager administrator password.

   Alternatively, instead of a password, type a pass-phrase that you were given by your vRealize Operations Manager administrator.

10 Click **Next**, and click **Finish**.

The administration interface appears, and it takes a moment for vRealize Operations Manager to finish adding the data node.

What to do next

After creating a data node, you have the following options.

- New, unstarted clusters:
  - Create and add more data nodes.
  - Create and add remote collector nodes.
  - Create a high availability master replica node.
  - Click **Start vRealize Operations Manager** to start the cluster, and log in to finish configuring the product.

    The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.
- Established, running clusters:
  - Create and add more data nodes.
  - Create and add remote collector nodes.
  - Create a high availability master replica node, which requires a cluster restart.

Advantages of an Expanding an Installation

A data node shares the load of performing vRealize Operations Manager analysis and it can also have an adapter installed to perform collection and data storage from the environment. You must have a master node before you add data nodes to form a cluster.

Resize your Cluster by Adding Nodes

You can deploy and configure additional nodes so that vRealize Operations Manager can support larger environments.

Figure 3-6. Workflow - Resize your cluster

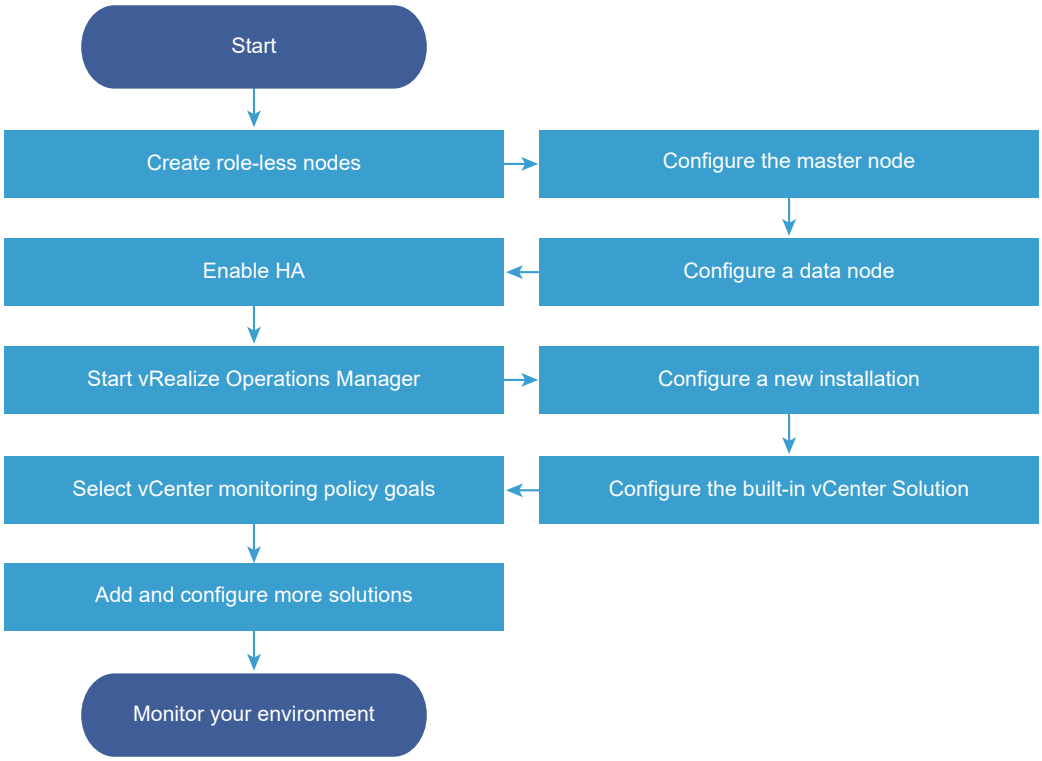

Gathering More Data by Adding a vRealize Operations Manager Remote Collector Node

You deploy and configure remote collector nodes so that vRealize Operations Manager can add to its inventory of objects to monitor without increasing the processing load on vRealize Operations Manager analytics.

Run the Setup Wizard to Create a Remote Collector Node

In distributed vRealize Operations Manager environments, remote collector nodes increase the inventory of objects that you can monitor without increasing the load on vRealize Operations Manager in terms of data storage, processing, or analysis.

Prerequisites

- Create nodes by deploying the vRealize Operations Manager vApp.

  During vApp deployment, select a remote collector size option.
- Ensure any remote adapter instance is running on the correct remote collector. If you have only one adapter instance, select Default collector group.
- Create and configure the master node.
- Note the fully qualified domain name (FQDN) or an IP address of the master node.

- Verify that there is one remote collector already added before you add another remote collector.

Note Remote collectors when added in parallel cause a cluster to crash.

Procedure

1. In a Web browser, navigate to the name or IP address of the deployed OVF that will become the remote collector node.

   The setup wizard appears, and you do not need to log in to vRealize Operations Manager.
2. Click **Expand an Existing Installation**.
3. Click **Next**.
4. Enter a name for the node, for example, **Remote–1**.
5. From the **Node Type** drop-down menu, select **Remote Collector**.
6. Enter the FQDN or IP address of the master node and click **Validate**.
7. Select **Accept this certificate** and click **Next**.

   If necessary, locate the certificate on the master node and verify the thumbprint.
8. Verify the vRealize Operations Manager administrator username of **admin**.
9. Enter the vRealize Operations Manager administrator password.

   Alternatively, instead of a password, type a passphrase that you were given by the vRealize Operations Manager administrator.
10. Click **Next**, and click **Finish**.

    The administration interface appears, and it takes several minutes for vRealize Operations Manager to finish adding the remote collector node.

What to do next

After creating a remote collector node, you have the following options.

- New, unstarted clusters:
  - Create and add data nodes.
  - Create and add more remote collector nodes.
  - Create a high availability master replica node.
  - Click **Start vRealize Operations Manager** to start the cluster, and log in to finish configuring the product.

    The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

- Established, running clusters:
  - Create and add data nodes.
  - Create and add more remote collector nodes.
  - Create a high availability master replica node, which requires a cluster restart.

Adding High Availability to vRealize Operations Manager

You can dedicate one vRealize Operations Manager cluster node to serve as a replica node for the vRealize Operations Manager master node.

Run the Setup Wizard to Add a Master Replica Node

You can convert a vRealize Operations Manager data node to a replica of the master node, which adds high availability (HA) for vRealize Operations Manager.

Note If the cluster is running, enabling HA restarts the cluster.

If you convert a data node that is already in use for data collection and analysis, adapters and data connections that were provided through that data node fail over to other data nodes.

You may add HA to the vRealize Operations Manager cluster at installation time or after vRealize Operations Manager is up and running. Adding HA at installation is less intrusive because the cluster has not yet started.

Prerequisites

- Create nodes by deploying the vRealize Operations Manager vApp.
- Create and configure the master node.
- Create and configure a data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the master node.

Procedure

1. In a Web browser, navigate to the master node administration interface.

   `https://master-node-name-or-ip-address/admin`
2. Enter the vRealize Operations Manager administrator username of **admin**.
3. Enter the vRealize Operations Manager administrator password and click **Log In**.
4. Under High Availability, click **Enable**.
5. Select a data node to serve as the replica for the master node.
6. Select the **Enable High Availability for this cluster** option, and click **OK**.

   If the cluster was online, the administration interface displays progress as vRealize Operations Manager configures, synchronizes, and rebalances the cluster for HA.

7 If the master node and replica node go offline, and the master remains offline for any reason while the replica goes online, the replica node does not take over the master role, take the entire cluster offline, including data nodes and log in to the replica node command line console as a root.

8 Open `$ALIVE_BASE/persistence/persistence.properties` in a text editor.

9 Locate and set the following properties:

```
db.role=MASTER
db.driver=/data/vcops/xdb/vcops.bootstrap
```

10 Save and close *persistence.properties*.

11 In the administration interface, bring the replica node online, and verify that it becomes the master node and bring the remaining cluster nodes online.

What to do next

After creating a master replica node, you have the following options.

- New, unstarted clusters:
  - Create and add data nodes.
  - Create and add remote collector nodes.
  - Click **Start vRealize Operations Manager** to start the cluster, and log in to finish configuring the product.

    The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.
- Established, running clusters:
  - Create and add data nodes.
  - Create and add remote collector nodes.

vRealize Operations Manager Cluster and Node Maintenance

You perform cluster and node maintenance procedures to help your vRealize Operations Manager perform more efficiently cluster and node maintenance involves activities such as changing the online or offline state of the cluster or individual nodes, enabling or disabling high availability (HA), reviewing statistics related to the installed adapters, and rebalancing the workload for better performance.

You perform most vRealize Operations Manager cluster and node maintenance using the Cluster Management page in the product interface, or the Cluster Status and Troubleshooting page in the administration interface. The administration interface provides more options than the product interface.

Table 3-3. Cluster and Node Maintenance Procedures

Procedure	Interface	Description
Change Cluster Status	Administration/Product	You can change the status of a node to online or offline. In a high availability (HA) cluster, taking the master or replica offline causes vRealize Operations Manager to run from the remaining node and for HA status to be degraded. Any manual or system action that restarts the cluster brings all vRealize Operations Manager nodes online, including any nodes that you had taken offline. If you take a data node that is part of a multi-node cluster offline and then bring it back online, the End Point Operations Management adapter does not automatically come back online. To bring the End Point Operations Management adapter online, select the End Point Operations Management adapter in the Inventory Explorer and click the **Start Collector** icon .
Enable or Disable High Availability	Administration	Enabling or disabling high availability requires the cluster to have at least one data node, with all nodes online or all offline. You cannot use Remote Collector nodes. Disabling high availability removes the replica node and restarts the vRealize Operations Manager cluster. After you disable high availability, the replica node vRealize Operations Manager converts back to a data node and restarts the cluster.
Generate Passphrase	Administration	You can generate a passphrase to use instead of the administrator credentials to add a node to this cluster. The passphrase is only valid for a single use.
Remove a Node	Administration	When you remove a node, you lose data that the node had collected unless you are running in high availability (HA) mode. HA protects against the removal or loss of one node. You must not re-add nodes to vRealize Operations Manager that you already removed. If your environment requires more nodes, add new nodes instead. When you perform maintenance and migration procedures, you should take the node offline, not remove the node.

Table 3-3. Cluster and Node Maintenance Procedures (continued)

Procedure	Interface	Description
Configure NTP	Product	The nodes in vRealize Operations Manager cluster synchronize with each other by standardizing on the master node time or by synchronizing with an external Network Time Protocol (NTP) source.
Rebalance the Cluster	Product	You can rebalance adapter, disk, memory, or network load across vRealize Operations Manager cluster nodes to increase the efficiency of your environment.

Cluster Management

vRealize Operations Manager includes a central page where you can monitor and manage the nodes in your vRealize Operations Manager cluster as well as the adapters that are installed on the nodes.

How Cluster Management Works

Cluster management lets you 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

In the left pane, select **Administration > Cluster Management**.

Cluster Management Options

The options include cluster-level monitoring and management features.

Table 3-4. Initial Setup Status Details

Option	Description
Cluster Status	Displays the online, offline, or unknown state of the vRealize Operations Manager cluster.
High Availability	Indicates whether HA is enabled, disabled, or degraded.

vRealize Operations Manager provides node-level information as well as a toolbar for taking nodes online or offline.

Table 3-5. 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. Master and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.

Table 3-5. Nodes in the vRealize Operations Manager Cluster (continued)

Option	Description
Cluster Role	Type of vRealize Operations Manager node: master, data, replica, or remote collector.
State	Running, Not Running, Going Online, Going Offline, Inaccessible, Failure, Error
Status	Online, offline, unknown, or other condition of the node.
Objects in Process	Total environment objects that the node currently monitors.
Metrics in Process	Total metrics that the node has collected since being added to the cluster.
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

In addition, there are adapter statistics for the selected node.

Table 3-6. 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 Being Collected	Total environment objects that the adapter currently monitors.
Metrics Being Collected	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 Post-Installation Considerations

After you install vRealize Operations Manager, there are post-installation tasks that might need your attention.

About Logging In to vRealize Operations Manager

Logging in to vRealize Operations Manager requires that you point a Web browser to the fully qualified domain name (FQDN) or IP address of a node in the vRealize Operations Manager cluster.

When you log in to vRealize Operations Manager, there are a few things to keep in mind.

- After initial configuration, the product interface URL is:

  https://*node-FQDN-or-IP-address*
- Before initial configuration, the product URL opens the administration interface instead.
- After initial configuration, the administration interface URL is:

  https://*node-FQDN-or-IP-address*/admin
- The administrator account name is admin. The account name cannot be changed.
- The admin account is different from the root account used to log in to the console, and does not need to have the same password.
- When logged in to the administration interface, avoid taking the node that you are logged into offline and shutting it down. Otherwise, the interface closes.
- The number of simultaneous login sessions before a performance decrease depends on factors such as the number of nodes in the analytics cluster, the size of those nodes, and the load that each user session expects to put on the system. Heavy users might engage in significant administrative activity, multiple simultaneous dashboards, cluster management tasks, and so on. Light users are more common and often require only one or two dashboards.

  The sizing spreadsheet for your version of vRealize Operations Manager contains further detail about simultaneous login support. See Knowledge Base article 2093783.
- You cannot log in to a vRealize Operations Manager interface with user accounts that are internal to vRealize Operations Manager, such as the maintenance Admin account.
- You cannot open the product interface from a remote collector node, but you can open the administration interface.
- For supported Web browsers, see the vRealize Operations Manager Release Notes for your version.

Secure the vRealize Operations Manager Console

After you install vRealize Operations Manager, you secure the console of each node in the cluster by logging in for the first time.

Procedure

1. Locate the node console in vCenter or by direct access. In vCenter, use Alt+F1 to access the login prompt.

   For security, vRealize Operations Manager remote terminal sessions are disabled by default.
2. Log in as **root**.

   vRealize Operations Manager prevents you from accessing the command prompt until you create a root password.

3 When prompted for a password, press Enter.

4 When prompted for the old password, press Enter.

5 When prompted for the new password, enter the root password that you want, and note it for future reference.

6 Re-enter the root password.

7 Log out of the console.

Log in to a Remote vRealize Operations Manager Console Session

As part of managing or maintaining the nodes in your vRealize Operations Manager cluster, you might need to log in to a vRealize Operations Manager node through a remote console.

For security, remote login is disabled in vRealize Operations Manager by default. To enable remote login, perform the following steps.

Procedure

1 Locate the node console in vCenter or by direct access. In vCenter, use Alt+F1 to access the login prompt.

2 Log in as **root**. If this is the first time logging in, you must set a root password.

   a When prompted for a password, press Enter.

   b When prompted for the old password, press Enter.

   c When prompted for the new password, enter the root password that you want, and note it for future reference.

   d Re-enter the root password.

3 To enable remote login, enter the following command:

```
service sshd start
```

About New vRealize Operations Manager Installations

A new vRealize Operations Manager installation requires that you deploy and configure nodes. Then, you add solutions for the kinds of objects to monitor and manage.

After you add solutions, you configure them in the product and add monitoring policies that gather the kind of data that you want.

Log In and Continue with a New Installation

To finish a new vRealize Operations Manager installation, you log in and complete a one-time process to license the product and configure solutions for the kinds of objects that you want to monitor.

Prerequisites

- Create the new cluster of vRealize Operations Manager nodes.

- Verify that the cluster has enough capacity to monitor your environment. See Sizing the vRealize Operations Manager Cluster.

Procedure

1. In a Web browser, navigate to the IP address or fully qualified domain name of the master node.
2. Enter the username **admin** and the password that you defined when you configured the master node, and click **Login**.

   Because this is the first time you are logging in, the administration interface appears.
3. To start the cluster, click **Start vRealize Operations Manager**.
4. Click **Yes**.

   The cluster might take from 10 to 30 minutes to start, depending on your environment. Do not make changes or perform any actions on cluster nodes while the cluster is starting.
5. When the cluster finishes starting and the product login page appears, enter the admin username and password again, and click **Login**.

   A one-time licensing wizard appears.
6. Click **Next**.
7. Read and accept the End User License Agreement, and click **Next**.
8. Enter your product key, or select the option to run vRealize Operations Manager in evaluation mode.

   Your level of product license determines what solutions you may install to monitor and manage objects.

   - Standard. vCenter only
   - Advanced. vCenter plus other infrastructure solutions
   - Enterprise. All solutions

   vRealize Operations Manager does not license managed objects in the same way that vSphere does, so there is no object count when you license the product.

   Note When you transition to the Standard edition, you no longer have the Advanced and Enterprise features. After the transition, delete any content that you created in the other versions to ensure that you comply with EULA and verify the license key which supports the Advanced and Enterprise features.
9. If you entered a product key, click **Validate License Key**.
10. Click **Next**.
11. Select whether or not to return usage statistics to VMware, and click **Next**.

12 Click **Finish**.

The one-time wizard finishes, and the vRealize Operations Manager interface appears.

What to do next

- Use the vRealize Operations Manager interface to configure the solutions that are included with the product.
- Use the vRealize Operations Manager interface to add more solutions.
- Use the vRealize Operations Manager interface to add monitoring policies.

Updating, Migrating and Restoring

You can update your existing vRealize Operations Manager deployments to a newly released version.

When you perform a software update, you need to make sure you use the correct PAK file for your cluster. A good practice is to take a snapshot of the cluster before you update the software, but you must remember to delete the snapshot once the update is complete.

If you have customized the content that vRealize Operations Manager provides such as alerts, symptoms, recommendations, and policies, and you want to install content updates, clone the content before performing the update. In this way, you can select the option to reset out-of-the-box content when you install the software update, and the update can provide new content without overwriting customized content.

Obtain the Software Update PAK File

Each type of cluster update requires a specific PAK file. Make sure you are using the correct one.

Download the Correct PAK files

To update your vRealize Operations Manager environment, you need to download the right PAK file for the clusters you wish to upgrade. Notice that only the Virtual Appliance clusters use an OS Update PAK file. Host name entries in the `/etc/hosts` of each node might be reset when applying the OS update PAK file for an update from vRealize Operations 6.0.x to version 6.1. You can manually update the hosts file after completing the software update.

Table 3-7. Specific PAK Files for Different Cluster Types

Cluster Type	OS Update	Product Update
Virtual Appliance clusters. Use both the OS and the product update PAK files.	`vRealize_Operations_Manager-VA-OS-xxx.pak`	`vRealize_Operations_Manager-VA-xxx.pak`
RHEL standalone clusters.		`vRealize_Operations_Manager-RHEL-xxx.pak`

Create a Snapshot as Part of an Update

It's a good practice to create a snapshot of each node in a cluster before you update a vRealize Operations Manager cluster. Once the update is complete, you must delete the snapshot to avoid performance degradation.

For more information about snapshots, see the vSphere Virtual Machine Administration documentation.

Procedure

1. Log into the vRealize Operations Manager Administrator interface at `https://<master-node-FQDN-or-IP-address>/admin`.
2. Click **Take Offline** under the cluster status.
3. When all nodes are offline, open the vSphere client.
4. Right-click a vRealize Operations Manager virtual machine.
5. Click **Snapshot** and then click **Take Snapshot**.
   a. Name the snapshot. Use a meaningful name such as "Pre-Update."
   b. Uncheck the **Snapshot the Virtual Machine Memory** check box.
   c. Uncheck the **Ensure Quiesce Guest File System (Needs VMware Tools installed)** check box.
   d. Click **OK**.
6. Repeat these steps for each node in the cluster.

What to do next

Start the update process as described in Install a Software Update.

How To Preserve Customized Content

When you upgrade vRealize Operations Manager, it is important that you upgrade the current versions of content types that allow you to alert on and monitor the objects in your environment. With upgraded alert definitions, symptom definitions, and recommendations, you can alert on the various states of objects in your environment and identify a wider range of problem types. With upgraded views, you can create dashboards and reports to easily identify and report on problems in your environment.

You might need to perform certain steps before you upgrade the alert definitions, symptom definitions, recommendations, and views in your vRealize Operations Manager environment.

- If you customized any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of vRealize Operations Manager, and you want to retain those customized versions, perform the steps in this procedure.

- If you did not customize any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of vRealize Operations Manager, you do not need to back them up first. Instead, you can start the upgrade, and during the upgrade select the check box named **Reset out-of-the-box content**.

Prerequisites

You previously customized versions of your alert definitions, symptom definitions, recommendations, or views.

Procedure

1. Before you begin the upgrade to vRealize Operations Manager, back up the changes to your alert definitions, symptom definitions, recommendations, and views by cloning them.
2. Start the upgrade of vRealize Operations Manager.
3. During the upgrade, select the check box named **Reset out-of-the-box content**.

Results

After the upgrade completes, you have preserved your customized versions of alert definitions, symptom definitions, recommendations, and views, and you have the current versions that were installed during the upgrade.

What to do next

Review the changes in the upgraded alert definitions, symptom definitions, recommendations, and views. Then, determine whether to keep your previously modified versions, or to use the upgraded versions.

Backup and Restore

Backup and restore your vRealize Operations Manager system on a regular basis to avoid downtime and data loss in case of a system failure. If your system does fail, you can restore the system to the last full or incremental backup.

You can backup and restore vRealize Operations Manager single or multi-node clusters by using vSphere Data Protection or other backup tools. You can perform full, differential, and incremental backups and restores of virtual machines.

To backup and restore vRealize Suite components by using vSphere Data Protection and NetBackup, see the Backup and Restore section in the vRealize Suite Information Center.

Note All nodes are backed up and restored at the same time. You cannot back up and restore individual nodes.

vRealize Operations Manager Software Updates

vRealize Operations Manager includes a central page where you can manage updates to the product software.

How Software Updates Work

The Software Update option lets you install updates to the vRealize Operations Manager product itself.

Where You Find Software Updates

Log in to the vRealize Operations Manager administration interface at https://*master-node-name-or-ip-address*/admin. On the left, click **Software Update**.

Software Update Options

The options include a wizard for locating the update PAK file and starting the installation, plus a list of updates and the vRealize Operations Manager cluster nodes on which they are installed.

Table 3-8. Software Update Options

Option	Description
Install a Software Update	Launch a wizard that allows you to locate, accept the license, and start the installation of a vRealize Operations Manager software update.
Node Name	Machine name of the node where the update is installed
Node IP Address	Internet protocol (IP) address of the node where the update is installed. Master and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.
Update Step	Software update progress in step x of y format
Status	Success, failure, in-progress, or unknown condition of the software update

Install a Software Update

If you have already installed vRealize Operations Manager, you can update your software when a newer version becomes available.

Note Installation might take several minutes or even a couple hours depending on the size and type of your clusters and nodes.

Prerequisites

- Create a snapshot of each node in your cluster. See Create a Snapshot as Part of an Update for details.
- Obtain the PAK file for your cluster. See Obtain the Software Update PAK File for details.
- 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**.

- The version 6.2.1 vRealize Operations Manager update operation has a validation process that identifies issues before you start to update your software. Although it is good practice to run the pre-update check and resolve any issues found, users who have environmental constraints can disable this validation check.

  To disable the pre-update validation check, perform the following steps:

  - Edit the update file to `/storage/db/pakRepoLocal/bypass_prechecks_vRealizeOperationsManagerEnterprise-buildnumberofupdate.json`.
  - Change the value to TRUE and run the update.

  Note If you disable the validation, you might encounter blocking failures during the update itself.

Procedure

1. Log into the master node vRealize Operations Manager Administrator interface of your cluster at `https://master-node-FQDN-or-IP-address/admin`.
2. Click **Software Update** in the left panel.
3. Click **Install a Software Update** in the main panel.
4. Follow the steps in the wizard to locate and install your PAK file.
   a. If you are updating a Virtual Appliance deployment, perform the OS update.

      This updates the OS on the virtual appliance and restarts each virtual machine.

   b. Install the product update PAK file.

      Wait for the software update to complete. When it does, the Administrator interface logs you out.
5. Read the **End User License Agreement** and **Update Information**, and click **Next**.
6. Click **Install** to complete the installation of software update.
7. Log back into the master node Administrator interface.

   The main Cluster Status page appears and cluster goes online automatically. The status page also displays the Bring Online button, but do not click it.
8. Clear the browser caches and if the browser page does not refresh automatically, refresh the page.

   The cluster status changes to Going Online. When the cluster status changes to Online, the upgrade is complete.

   Note If a cluster fails and the status changes to offline during the installation process of a PAK file update then some nodes become unavailable. To fix this, you can access the Administrator interface and manually take the cluster offline and click **Finish Installation** to continue the installation process.

9 Click **Software Update** to check that the update is done.

A message indicating that the update completed successfully appears in the main pane.

What to do next

Delete the snapshots you made before the software update.

Note Multiple snapshots can degrade performance, so delete your pre-update snapshots after the software update completes.

Install a vRealize Operations Manager Software Update from the Administration Interface

You activate the vRealize Operations Manager product or its additional solutions by registering licenses.

Prerequisites

- Know the name and location of the software update 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 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 username and password for the master node.

3 On the left, click **Software Update**.

4 Click **Install a Software Update**.

5 Follow the wizard to locate and install your copy of *update-filename*.pak.

Installation completes in a couple of minutes, and the administrator interface logs you out. If you are not logged out automatically after 5 minutes, refresh the page in your browser.

6 Log back in to the master node administrator interface, and click **Software Update** again.

7 Verify that update name appears on the right. If the update does not appear, wait a few minutes, and refresh the page in your browser.

Migrate a vCenter Operations Manager Deployment into this Version

By importing data, an established or production version of vRealize Operations Manager can assume the monitoring of a vCenter Operations Manager deployment.

You cannot migrate vCenter Operations Manager directly to this version of vRealize Operations Manager. Instead, you follow a two-step process:

1 Migrate and import vCenter Operations Manager 5.8.x into vRealize Operations Manager 6.0.x as described in the version 6.0.x documentation.

2 Use the vRealize Operations Manager **Software Update** option to update vRealize Operations Manager 6.0.x to this version.

Note Make sure your vCenter Operations Manager 5.8.x and vRealize Operations Manager 6.0.x instances are on the same physical network. Otherwise the data import may not work. Data import process fails when source (vCenter Operations Manager 5.x) is separated from the destination vRealize Operations Manager 6.x environment by a slow network connection (WAN). Data import over a connection that is slower than LAN speed is not supported. For more information, see the Knowledge Base article 2141964.

Configuring

4

You configure objects, alerts, actions, policies, dashboards, and reports, in vRealize Operations Manager to effectively monitor your environment. You use administration settings to manage your environment.

Configure solutions in vRealize Operations Manager to connect to and analyze data from external data sources in your environment. Once connected, you use vRealize Operations Manager to monitor and manage objects in your environment. Solutions that are installed together with vRealize Operations Manager include vSphere, End Point Operations, Log Insight, vRealize Automation, VMware vSAN, and Business Management. Configure these adapters to connect to and integrate with these instances.

Create alert definitions so that whenever there is a problem, vRealize Operations Manager triggers alerts and provides recommendations to resolve the problem. The process of configuring alerts involves defining alerts, symptoms, and recommendations.

Enable actions to address a problem in the monitored environment. The actions let you resolve a problem by remaining in the vRealize Operations Manager environment itself.

Create a policy to define rules for vRealize Operations Manager to use. You can use a policy to analyze and display information about the objects in your environment.

Define compliance standards to determine the compliance of your objects. You can use vRealize Operations Manager alert definitions to create compliance standards that notify you when an object does not comply with a required standard.

Create super metrics to give you a big picture of your environment. A super metric is a mathematical formula that contains one or more metrics. It is a custom metric that you design and is useful when you need to track combinations of metrics, either from a single object or from multiple objects. If a single metric cannot tell you what you need to know about the behavior of your environment, you can define a super metric.

Create 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.

Create views to interpret metrics, properties, and policies of various monitored objects including alerts. Generate a report to capture details related to current or predicted resource needs. A report is a scheduled snapshot of views and dashboards.

This chapter includes the following topics:

- Connecting vRealize Operations Manager to Data Sources
- Configuring Alerts and Actions
- Configuring Policies
- Configuring Compliance
- Configuring Super Metrics
- Configuring Objects
- Configuring Data Display
- Configuring Administration Settings
- About the vRealize Operations Manager Administration Interface
- Configuring and Using Workload Balance

Connecting vRealize Operations Manager to Data Sources

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

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

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

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

VMware vSphere Solution in vRealize Operations Manager

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

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

Configuring the vSphere Solution

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

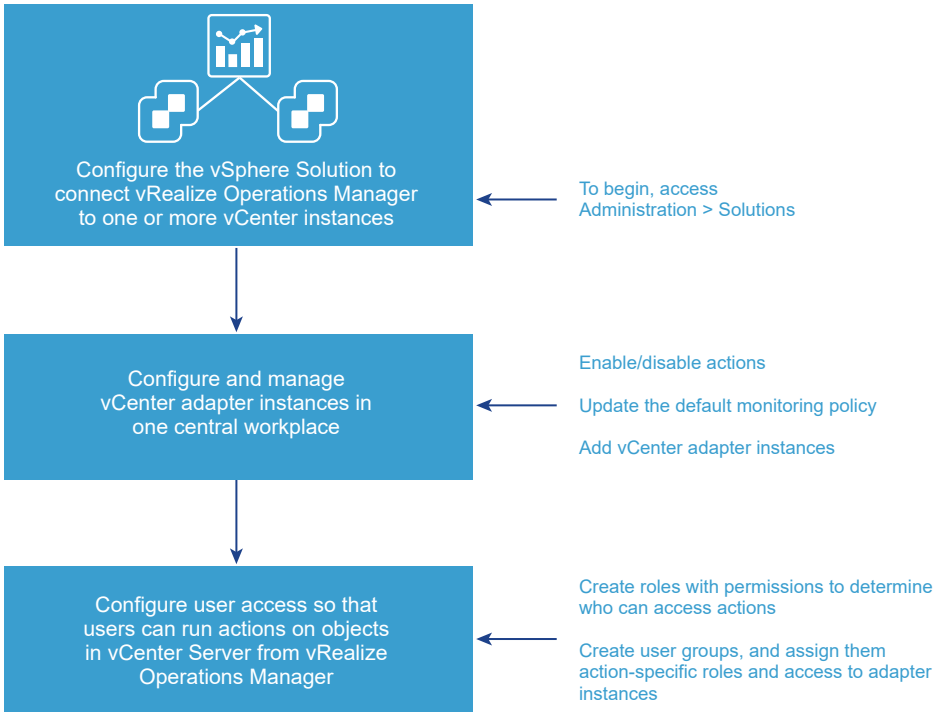

How Adapter Credentials Work

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

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

Controlling User Access to Actions

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

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

Configure a vCenter Adapter Instance in vRealize Operations Manager

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

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

Prerequisites

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

Procedure

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

   The vCenter Server FQDN or IP address must be reachable from all nodes in the vRealize Operations Manager cluster.
5. To add credentials for the vCenter Server instance, click the **Add** icon, and enter the required credentials.

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

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

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

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

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

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

   For information about these advanced settings, see the Manage Solution - VMware vSphere Solution Workspace Options.

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

   For information about monitoring goals, see the Manage Solution - VMware vSphere Solution Workspace Options.

11 Click **Save Settings**.

   The adapter instance is added to the list.

Results

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

For information about the network port that vRealize Operations Manager uses to communicate with a vCenter Server system and vRealize Operations Manager components, see Port Requirements for vRealize Operations Manager.

What to do next

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

Configure User Access for Actions

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

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

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

Procedure

1 On the menu, click **Administration** and in the left pane click **Access > Access Control**.

2 To create a role:

   a Click the **Roles** tab.

   b Click the **Add** icon, and enter a name and description for the role.

3 To apply permissions to the role, select the role, and in the Permissions pane, click the **Edit** icon.

   a Expand **Environment**, and then expand **Action**.

   b Select one or more of the actions, and click **Update**.

4 To create a user group:

   a Click the **User Groups** tab, and click the **Add** icon.

   b Enter a name for the group and a description, and click **Next**.

   c Assign users to the group, and click the **Objects** tab.

   d Select a role that has been created with permissions to run actions, and select the **Assign this role to the user** check box.

   e Configure the object privileges by selecting each adapter instance to which the group needs access to run actions.

   f Click **Finish**.

What to do next

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

Manage Solution - VMware vSphere Solution Workspace Options

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

Where You Find the Manage Solution - VMware vSphere Workspace

On the menu, click **Administration** and in the left pane click **Solutions**. On the **Solutions** tab, select **VMware vSphere** and click the **Configure** icon on the toolbar.

Manage Solution - VMware vSphere Workspace Options

Configure and modify adapter instances, and define monitoring goals on the Manage Solution page.

Table 4-1. Manage Solution Page Options

Option	Description
Adapter Type list	Provides a list of the adapters included in the solution. Configured adapters provide the settings and credentials that vRealize Operations Manager must communicate with your vCenter Server instances or action instances. 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. To reset the default content, navigate to the Solutions configuration page, and click **Administration > Solutions**. Click the VMware vSphere solution, click **Configure**, and in the Manage Solution workspace, click **Reset Default Content**. The option named **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 menu item to overwrite alert definitions and symptom definitions. If you do not overwrite alert and symptom definitions, compliance rules will use a mixture of new and outdated definitions.
Instance Name list	List of configured adapter instances based on the selected adapter type. This list is blank until you configure at least one instance.
Instance Settings	Settings used to identify the target vCenter Server instance. ■ 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.
Basic Settings	Minimum settings used to connect to the target vCenter Server. ■ vCenter Server. Enter the FQDN or IP address of the target vCenter Server instance. The FQDN or IP address must be reachable from all nodes in the vRealize Operations Manager cluster. ■ Credentials. Click the **Add** icon to add credential details.
vCenter Actions	Settings used to configure the adapter to run actions on objects in the vCenter Server from vRealize Operations Manager, ■ Enable Actions? The vCenter adapter is configured to run actions on objects in the vCenter Server instance by default. Select **Disable** if you do not want the adapter to run actions. Select **Enable** to run actions on objects. ■ (Optional) Alternate Action Credentials. You can use the same credentials you provided to connect to the vCenter Server to run actions, or click this menu item to provide alternative credentials. ■ Test Connection. Click to verify that the provided credentials can connect to the target vCenter Server and so that you can validate the certificate. The certificate presented is the leaf certificate for the vCenter Server instance, not the complete certificate chain. Click **OK** only if the certificate presented in the dialog box matches the certificate for your target vCenter Server.

Table 4-1. Manage Solution Page Options (continued)

Option	Description
Advanced Settings	Provides options related to designating specific collectors to manage this adapter instance, managing object discovery and change events.
Collectors/Groups	Determines which vRealize Operations Manager collector is used to manage the adapter processes. If you have only one adapter instance, select **Default collector group**. If you have multiple collectors in your environment, and you want to distribute the workload to optimize performance, select the collector to manage the adapter processes for this instance.
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 adapter. ■ If the value is true, vRealize Operations Manager collects 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 the objects that are present on the target system when you configure the adapter instance.
Process Change Events	Determines whether the adapter 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.
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.

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 new default policy.

Table 4-2. 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.
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 assess and operate your vSphere objects. When you enable these alerts, vRealize Operations Manager assesses your objects against the *vSphere Hardening Guide* rule.

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.

End Point Operations Management Solution in vRealize Operations Manager

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

End Point Operations Management Agent Installation and Deployment

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

Prepare to Install the End Point Operations Management Agent

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

Prerequisites

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

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

- Define the agent `HQ_JAVA_HOME` location.

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

Supported Operating Systems for the End Point Operations Management Agent

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

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

Table 4-3. 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.

Note that 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, VMware recommends 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 4-4. 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 4-5. Additional Silent Command Line Installer Parameters

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

Install the Agent on an AIX Platform

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

Prerequisites

1 Install IBM Java 7.

2 Add the latest JCE from the IBM JRE security directory: `JAVA_INSTALLATION_DIR/jre/lib/security`. For more information, see Downloading and installing the unrestricted JCE policy files

Procedure

1 When you configure the PATH variable, add `/usr/java7_64/jre/bin:/usr/java7_64/bin` or *PATH*=`/usr/java7_64/jre/bin:/usr/java7_64/bin:$PATH`.

2 Configure `HQ_JAVA_HOME=path_to_current_java_directory`.

For more information on setting up and checking your AIX environment, see https://www.ibm.com/support/knowledgecenter/SSYKE2_7.0.0/com.ibm.java.aix.70.doc/diag/problem_determination/aix_setup.html.

3 Download the noJre version of the End Point Operations Management agent and install the agent on an AIX machine.

4 For agent installation information, see Install the Agent on a Linux Platform from an Archive

Install the Agent on a Solaris Platform

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

Prerequisites

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

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

Procedure

1 When you configure the PATH variable, add `/usr/java7_64/jre/bin:/usr/java7_64/bin` or *PATH*=`/usr/java7_64/jre/bin:/usr/java7_64/bin:$PATH`.

2 Configure `HQ_JAVA_HOME=path_to_current_java_directory`.

3 Download and install the noJre version of the End Point Operations Management agent on a Solaris machine.

4 For agent installation information, see Install the Agent on a Linux Platform from an Archive

Java Prerequisites for the End Point Operations Management Agent

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

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

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

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

Configuring JRE Locations for End Point Operations Management Components

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

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

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

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

How the Agent Resolves its JRE

The agent resolves its JRE based on platform type.

UNIX-like Platforms

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

1. `HQ_JAVA_HOME` environment variable
2. Embedded JRE
3. `JAVA_HOME` environment variable

Linux Platforms

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

Windows Platforms

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

1. `HQ_JAVA_HOME` environment variable

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

2 Embedded JRE

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

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

   The End Point Operations Management agent configuration dialog 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 prompts you to provide the address and port of the vRealize Operations Manager server, and other connection-related data.

5. Overriding Agent Configuration Properties

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

6 End Point Operations Management Agent Properties

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

What to do next

Start the End Point Operations Management agent.

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

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

Procedure

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

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

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

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

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

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

Specify the End Point Operations Management Agent Setup Properties

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

Agent-server setup requires a minimum set of properties.

Procedure

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

Property	Property Definition
agent.setup.serverIP	Specify the address or hostname of the vRealize Operations Manager server.
agent.setup.serverSSLPort	The default value is the standard SSL vRealize Operations Manager server listen port. If your server is configured for a different listen port, specify the port number.
agent.setup.serverLogin	Specify the user name for the agent to use when connecting to the vRealize Operations Managerserver. 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 user name specified in `agent.setup.camLogin`, 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

The End Point Operations Management agent configuration dialog 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 prompts you to provide the address and port of the vRealize Operations Manager server, and other connection-related data.

The agent configuration dialog 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.

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, 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 anEnd 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
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.

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 would 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.

Option	Description
status	Queries the status of the agent's JVM process.
dump	Runs a thread dump for the agent process, and writes the result to the `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, you must delete the original machine from vRealize Operations Manager to enable the new operating system relationship with the virtual machine to be created.

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.

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 automatically 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 a 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 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, and 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 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 4-6. ssl Option

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

Table 4-7. 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 4-8. port Option

Option Information	Value
Description	Port
Default	A default value for port is usually 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 4-9. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	true
Type	int

Table 4-9. sotimeout Option (continued)

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

Table 4-10. 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 4-11. 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 4-12. 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 4-13. 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 re-directed. 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 4-14. pattern Option

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

Table 4-15. 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 4-16. 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 4-17. Credential Option

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

ICMP Configuration Options

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

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

Table 4-18. 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 4-19. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A

Table 4-19. sotimeout Option (continued)

Option Information	Value
Type	int
Notes	The maximum period of time 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 4-20. port Option

Option Information	Value
Description	Port
Default	A default value for port is usually 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 4-21. 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 you use the IP address of the machine on which the remote check is to run, not the host name.

Table 4-22. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A
Type	int

Table 4-22. sotimeout Option (continued)

Option Information	Value
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 Managementagent 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 labelled 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
```

Log Insight

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

Log Insight Page

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

To access theLog 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.

Configuring vRealize Log Insight with vRealize Operations Manager

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

Configuring the vRealize Log Insight Adapter in vRealize Operations Manager

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

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

Prerequisites

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

Procedure

1. In the menu, click **Administration**, and then in the left pane click **Solutions**.
2. From the Solutions page, click VMware vRealize Log Insight.
3. Click the **Configure** icon. You see the Manage Solution-VMware vRealize Log Insight dialog box.
4. In the Manage Solutions dialog box perform the following steps:
   - Enter a name in the **Display Name** text box.
   - Enter the IP address in the **Log Insight server** text box of the vRealize Log Insight you have installed and want to integrate with.
   - Click **Test Connection** to verify that the connection is successful.
   - Click **Save Settings**.

- Click **Close**.

5 From the vRealize Operations Manager Home page, click **Log Insight** from the left pane. If you see a statement at the bottom of the page, click the link and accept the certificate exception in vRealize Log Insight or contact your IT support for more information.

6 From the vRealize Operations Manager Home page, click **Log Insight** from the left pane and enter the user name and password of the vRealize Log Insight instance you have installed.

Configuring vRealize Operations Manager in vRealize Log Insight

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

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

Prerequisites

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

Procedure

1 From the Administration page of vRealize Log Insight, click the **vRealize Operations** icon from the left pane. You see the vRealize Operations Integration pane.

2 In the **Hostname** and **Username** text boxes, enter the IP address and hostname of the vRealize Operations Manager instance you want to integrate with.

3 In the **Password** text box, select **Update Password** and enter the password of the vRealize Operations Manager instance you want to integrate with.

4 Select the **Enable launch in context** option.

5 Click **Test Connection** to verify that the connection is successful.

6 Click **Save**.

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

Business Management

When vRealize Operations Manager is integrated with vRealize Business for Cloud, you can display infrastructure performance and cost information in the Business Management page.

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

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

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

Configure the vRealize Business for Cloud Adapter

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

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

Procedure

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

   If you have one adapter instance, select **Default collector group**. If you have multiple collectors in your environment, to distribute the workload and optimize performance, select the collector to manage the adapter processes for this instance.
7. Click **Save Settings** to finish configuration of the adapter, and click **Close**.

What to do next

To view data in the Business Management page:

1. On the menu, click **Administration** and in the left pane click **Business Management**.
2. Click the link at the bottom of the Business Management page and log into vRealize Business for Cloud.
3. Accept the certificate exception.

vRealize Automation Solution

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

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

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

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

Supported vRealize Automation Versions

The vRealize Automation solution is supported with vRealize Automation 7.x versions.

If you upgrade from a previous version to vRealize Operations Manager 6.6, that has the vRealize Automation Management Pack 2.1 installed and vRealize Automation 6.2.5 registered as an adapter instance, the following behavior is observed:

- vRealize Automation Management Pack 2.1 is upgraded to 3.0.
- The collection state fails for the vRealize Automation 6.2.5 adapter instance and the collection status is `None`.

Object Types and Relationships

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

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

- Tenant
- Reservation
- Business Group
- Deployment (appears only when monitoring vRealize Automation 7.0 and later environments)
- Blueprint
- Managed Resources

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.

- The tenant object type has a parent-child relationship with business group and reservation object types. The business group and reservation object types interact with each other.

- The tenant, reservation, deployment, and business group object types have a parent-child relationship with the virtual machines and VMware vCenter adapter objects such as clusters.

The fabric group, network profile, reservation policy, and storage reservation policy object types do not collect data.

vRealize Automation Workload Placement

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

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

Procedure

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

Port Information

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

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

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

Security Guidelines

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

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

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

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

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

Configuring vRealize Automation

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

Prerequisites

- The super user must have the following privileges:
  - Infrastructure administrator rights for all tenants.
  - Tenant administrator rights for all tenants.
  - Software architect roles for all tenants (required only for vRealize Automation 7.0 and later installations).
- Configure the vCenter adapter instance that is added by the vRealize Automation solution, for workload placements of the clusters managed by vRealize Automation.
- Use only DNS names and not IP addresses when you configure the vRealize Automation solution in a vRealize Automation distributed setup. Add host file entries on all vRealize Operations Manager nodes in the `/etc/hosts` location if the DNS is not reachable using vRealize Operations Manager.

Procedure

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

3 Configure the solution.

Option	Description
Display Name	The name for the adapter instance.
Description	(Optional) The description of the adapter instance.
vRealize Automation Appliance URL	The URL of the vRealize Automation CAFÉ appliance from which you are collecting data. Specify 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 HostName 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 on 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 note below. ■ **SuperUser Password**. The password of the vRealize Automation super user.
Advanced Settings	To configure the advanced settings, click on the drop-down menu.
Collectors/Groups	The collector on which the vRealize Automation solution 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.
Tenants	Collects data for specific tenants associated with vRealize Automation. Configure the tenants in the following manner to collect data: ■ * (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 the vRealize Automation solution continues to collect data for the other tenants. Failure in collecting data for a tenant that does not have the required privileges is logged in the `adapter.log` file. ■ **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 along with the compute clusters under managed resources. ■ **Disabled**: Collects and monitors data for only the reservation object type along with the compute clusters under managed resources.

Option	Description
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).
Autodiscovery	Discover objects automatically. ■ To set automatic discovery for objects, select **True**. ■ To set the automatic discovery off, select **False**.

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

If one of the tenant connections is successful, Test Connection is successful.

5 Click **Save Settings**.

Set the Collection Time Interval

The recommended collection time interval is configured in the `vcac.properties` file.

The default time is as follows:

```
# Timing in minutes.
defaultAdapterCollectionTime=15
# 4 hours
defaultTenantCollectionTime=240
#1 hour
defaultBusinessGroupCollectionTime=60
defaultBlueprintCollectionTime=60
```

You can change the recommended collection time interval in the `vcac.properties` file. You can access the `vcac.properties` file from: `/usr/lib/vmware-vcops/user/plugins/inbound/vCACAdapter3/conf/vcac.properties`. Restart the vRealize Automation adapter instance after you make the changes to the `vcac.properties` file.

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 4-23. 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 may be increased or decreased based on speed or performance requirements.	2
querySuiteAPIPageSize	The number of the items to fetch in a suite API call.	100
queryVraAPIPageSize	The number of the items to fetch in a single CAFE query.	100 **Note** It is recommended that you keep the maximum value as 100. Refer to the sizing guidelines for large scale environment guidelines: Sizing Guidelines

vSAN

You can make vSAN operational in a production environment by using dashboards to evaluate, manage, and optimize the performance of vSAN objects and vSAN-enabled objects in your vCenter Server system.

vSAN extends the following features:

- Discovers vSAN disk groups in a vSAN datastore.
- Identifies the vSAN-enabled cluster compute resource, host system, and datastore objects in a vCenter Server system.
- Automatically adds related vCenter Server components that are in the monitoring state.

Configure a vSAN Adapter Instance

When configuring an adapter instance for vSAN, you add credentials for a vCenter Server.

Prerequisites

Only vCenter Server systems that are configured for both the vCenter adapter and the vSAN adapter appear in the inventory tree under the vSAN and Storage Devices. Verify that the vCenter Server that you use to configure the vSAN adapter instance is also configured as a vCenter adapter instance for the VMware vSphere® solution. If not, add a vCenter adapter instance for that vCenter Server.

Procedure

1. In the vCenter Server text box, enter the FQDN or IP address of the vCenter Server instance to which you are connecting.

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

2. To add credentials on the Manage Solution page, click the plus sign.

   a. In the Credential name text box, enter the name by which you are identifying the configured credentials.

   b. Type the User name and Password for the vCenter Server instance.

   c. Click **OK**.

   You configured credentials to connect to a vCenter Server instance.

3. Click **Advanced Settings**.

4. For Collection Interval, select a value of five minutes or greater.

   The vSAN adapter collects Health Check Service and Performance Service metrics from vSphere objects. The Health Check Service interval is configured in the vSphere interface and is 60 minutes by default. If the Health Check Service interval is 60 minutes and the vSAN adapter collection interval is 5 minutes, the vSAN adapter reports:

   - Same Health Check Service metrics 12 times.
   - Most recent Performance Service metrics for each cycle.

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

6. Accept the vCenter Server security certificate.

7. Click **Save Settings**.

Results

The adapter is added to the Adapter Instance list and is active.

What to do next

To verify that the adapter is configured and collecting data from vSAN objects, wait a few collection cycles, then view application-related data.

- Inventory Explorer. Verify that all the objects related to the vSAN instance are listed. Objects should be in the collecting state and receiving data.
- Dashboards. Verify that vSAN Capacity Overview, Optimize vSAN Deployments, 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-enabled Datastores
- vSAN Disk Groups
- vSAN Datastore related VMs
- vSAN Witness Hosts (optional)

Verify that the Adapter Instance is Connected and Collecting Data

You configured an adapter instance of vSAN with credentials for a vCenter Server. Now you want to verify that your adapter instance can retrieve information from vSAN objects in your environment.

To view the object types, in the menu, click **Administration > Configuration > Inventory Explorer > Adapter Instances > vSAN Adapter Instance > <*User_Created_Instance*>**.

Table 4-24. Object Types that vSAN Discovers

Object Type	Description
vSAN Adapter Instance	The vRealize Operations Management Pack for vSAN instance.
vSAN Cluster	vSAN clusters in your data center.
vSAN Datastore	vSAN datastores in your data center.
vSAN Disk Group	A collection of SSDs and magnetic disks used by vSAN.
vSAN Fault Domain	A tag for a fault domain in your data center.
vSAN Host	vSAN hosts in your data center.
vSAN Witness Host	A tag for a witness host of a stretched cluster, if the stretched cluster feature is enabled on the vSAN cluster.
vSAN World	A vSAN World is a group parent resource for all vSAN adapter instances. vSAN World displays aggregated data of all adapter instances and a single root object of the entire vSAN hierarchy.
Cache Disk	A local physical device on a host used for storing VM files in vSAN.
Capacity Disk	A local physical device on a host used for read or write caching in vSAN

The vSAN adapter also monitors the following objects discovered by the VMware vSphere adapter.

- Cluster Compute Resources
- Host System
- Datastore

Procedure

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

2 In the list of tags, expand **Adapter Instances** and expand **vSAN Adapter Instance**.

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

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
      ]
   }
```

Installing Optional Solutions in vRealize Operations Manager

You can extend the monitoring capabilities of vRealize Operations Manager by installing optional solutions from VMware or third parties.

VMware solutions include adapters for Storage Devices, Log Insight, NSX for vSphere, Network Devices, and VCM. Third-party solutions include AWS, SCOM, EMC Smarts, and many others. To download software and documentation for optional solutions, visit the VMware Solution Exchange at https://marketplace.vmware.com/vsx/.

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

If you upgrade from an earlier version of vRealize Operations Manager, your management pack files are copied to the `/usr/lib/vmware-vcops/user/plugins/.backup` file in a folder with the date and time as the folder name. Before migrating your data to your new vRealize Operations Manager instance, you must configure the adapter instances again. If you have customized the adapter, your adapter customizations are not included in the migration, and you must reconfigure the customizations.

If you update a management pack in vRealize Operations Manager to a newer version, and you have customized the adapter, your adapter customizations are not included in the upgrade, and you must reconfigure them.

Solutions in vRealize Operations Manager

vRealize Operations Manager includes a page where you can add and manage solutions, which include the adapters that connect to the data to monitor and manage.

How Solutions Work

Solutions can include content and adapters. vRealize Operations Manager uses adapters to manage communication and integration with other products, applications, and functions.

Where You Find Solutions

In the menu, click **Administration** and in the left pane click **Solutions**.

Data Collection Notifications

The **Data Collection** 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.

Solutions Options

The solutions list includes a toolbar of options.

Table 4-25. Solutions Toolbar Options

Option	Description
Add	Start a wizard to find, upload, license, and install a solution PAK file.
Configure	Open a window in which you control settings such as network addresses or credentials that allow the solution to connect to data. Configuration varies by solution.
Remove solution	Click to remove the selected solution. If the solution is designed for vRealize Operations Manager 6.6, the adapter and all its associated data is removed. For solutions that have not been designed for vRealize Operations Manager 6.6, all deployed adapter instances, metadata, and data including alert and symptom definitions is removed. Content such as the solution dashboards, views, reports, and metric configurations is not removed.
Show	Filter the list of solutions to show configured, unconfigured, or all solutions.

The solutions data grid is a list of solutions that were added. You must configure solution components so that vRealize Operations Manager can collect data.

Table 4-26. Solutions Data Grid Options

Option	Description
Name	Name that the vendor or manufacturer gave to the solution.
Description	Typically, an indication of what the solution monitors or what data source its adapter connects to.
Version	Version and build number identifiers of the solution.
Provided By	Vendor or manufacturer that created the solution.
Licensing	Indicates that the solution requires a license.
Adapter Status	Indicates the status of the solution. Data receiving shows that the solution is collecting data.

The details area includes a toolbar of options.

Table 4-27. Solution Details Toolbar Options

Option	Description
Start Collecting	Turn on data collection through the selected adapter.
Stop Collecting	Do not collect data through the selected adapter.
Reload	Refresh the list of details.

The details data grid displays additional information for the selected solution.

Table 4-28. Solution Details Data Grid Options

Option	Description
Adapter Type	Name that the vendor or manufacturer gave to the adapter.
Adapter Instance Name	Name that the installing user gave to this unique installation of the adapter.
Credential Name	Name that the installing user gave to the set of login credentials used to connect to the data source.
Collector	Indicates where vRealize Operations Manager is receiving the collected data. Typically, the name combines the adapter and the vRealize Operations Manager node names.
Collection State	Indicates whether the adapter is enabled for data collection.
Collection Status	Indicates whether the adapter has collected any data.

Add Solutions Wizard

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 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**. Select the solution you want to install, and click the **Add** icon.

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. 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**.

Table 4-29. 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 out-of-the-box 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** Clicking **Next** installs the solution.
Page 3	
Installation Details	Review the installation progress, including the vRealize Operations Manager nodes where the adapter was installed.

Manage Solutions Workspace

Solutions include adapters that you must configure so that vRealize Operations Manager can collect data from or send data to the target system.

You can configure adapters associated with solutions that are provided with or that you add to vRealize Operations Manager. After you have configured the adapter, vRealize Operations Manager can communicate with the target system. You can access the Manage Solutions workspace at any time to modify your adapter configurations.

Where You Manage Solutions

On the menu, click **Administration** and in the left pane click **Solutions**. Select the solution you want to manage.

The options available depend on the selected solution.

Manage the vSphere Solution

To view the manage solution workspace options of the vSphere solution, see Manage Solution - VMware vSphere Solution Workspace Options.

Managing Solution Credentials

Credentials are the user accounts that vRealize Operations Manager uses to enable one or more solutions and associated adapters, and to establish communication with the target data sources. The credentials are supplied when you configure each adapter. You can add or modify the credential settings outside the adapter configuration process to accommodate changes to your environment.

For example, if you are modifying credentials to accommodate changes based on your password policy, the adapters configured with these credentials begin using the new user name and password to communicate between vRealize Operations Manager and the target system.

Another use of credential management is to remove misconfigured credentials. If you delete valid credentials that were in active use by an adapter, you disable the communication between the two systems.

If you need to change the configured credential to accommodate changes in your environment, you can edit the credential settings without being required to configure a new adapter instance for the target system. You can edit credential settings by clicking **Administration** on the menu, and then clicking **Credentials**.

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

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 4-30. Credentials Options

Option	Description
Toolbar options	Manages the selected credential. ■ Add New Credentials. Add new credentials for an adapter type that you can later apply when configuring an adapter. ■ Edit Selected Credentials. 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 Selected Credential. Deletes 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.

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

Table 4-31. 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 4-32. Control Group Summary Grid

Options	Description
Collector Group toolbar	To manage collector groups, use the toolbar icons. ■ Add. Add a collector group ■ Edit. Modify the collector group by adding or removing remote collectors. ■ Delete. Remove the selected collector group. ■ Rebalance collector group. If you have permissions to manage clusters, you can rebalance the workload across the collectors and the remote collectors in the collector group. You can only rebalance one collector group at a time. 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 4-33. 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

Configuring Alerts and Actions

In VMware vRealize Operations Manager, alerts and actions play key roles in monitoring the objects.

All Alerts

The All 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 All 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 filter the columns in the data grid, click the small box on the lower left of the alert list.

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 links to additional information. A **Run Action** button may appear in the details. Hover over the button to learn what recommendation is performed if you click the button. Click the X at the top right of the alert details to return to the list view.

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**.

All 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 4-34. 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.
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.
Go to Alert Definition	Switches to the Alert Definitions page, with the definition for the previously selected alert displayed.
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 4-35. 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 "All Alerts Data Grid Options" table, below.

Table 4-35. Group By Options (continued)

Option	Description
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 4-36. 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.
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.

Table 4-36. All Filters (continued)

All Filters	Descriptions
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.

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 4-37. All 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.

Table 4-37. All Alerts Data Grid (continued)

Option	Description
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.

Types of Alerts

Different types of alerts are triggered on a certain object.

The alerts are of three types:

- Health Alerts
- Risk Alerts
- Efficiency Alerts

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.

How Health Alerts Work

All the health alerts generated for you managed objects appear in the list.

You can manage the alerts in the list using the toolbar options, click the alert name to see the alert details for the affected object, or click the name of the object on which the alert was generated to see the object details.

Health Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to cancel, suspend, or manage ownership. You can select multiple rows in the list using Shift+click, Control +click. Use the data grid to view the alerts. You can click the alert name to view the alert details or object name to view the object details.

Table 4-38. Health Alerts Toolbar Options

Option	Description
Open in external application	Actions you can run on the selected object. For example, Open Virtual Machine in vSphere Client.
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.
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.
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.

The Health Alerts data grid provides a list of generated alerts that you use to resolve problems in your environment.

Table 4-39. Health Alerts Data Grid Options

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**. The possible values include: ■ Critical ■ Immediate ■ Warning ■ Information By default, alerts are sorted by criticality. Presorting the alerts list by criticality displays critical alerts at the top of the list. If you change the sort order, the sort is saved with your preferences in the global alerts list, and the Health, Risk, and Efficiency alerts lists.
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.
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.
Status	Current state of the alert. Possible values include Active or Canceled.
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.

Table 4-39. Health Alerts Data Grid Options (continued)

Option	Description
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.
Owner	Name of the user who owns the alert.
Created On	Date and time when 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.

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.

How Risk Alerts Work

All the risk alerts generated for your managed objects appear in the list.

You can manage the alerts in the list using the toolbar options, click the alert name to see the alert details for the affected object, or click the name of the object on which the alert was generated to see the object details.

Risk Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to cancel, suspend, or manage ownership. You can select multiple rows in the list using Shift+click, Control +click. Use the data grid to view the alerts. You can click the alert name to view the alert details or object name to view the object details.

Table 4-40. Risk Alerts Toolbar Options

Option	Description
Open in external application	Actions you can run on the selected object. For example, Open Virtual Machine in vSphere Client.
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.
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.
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.

The Risk Alerts data grid provides a list of generated alerts that you use to resolve problems in your environment.

Table 4-41. Risk Alerts Data Grid Options

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**. The possible values include: ■ Critical ■ Immediate ■ Warning ■ Information By default, alerts are sorted by criticality. Presorting the alerts list by criticality displays critical alerts at the top of the list. If you change the sort order, the sort is saved with your preferences in the global alerts list, and the Health, Risk, and Efficiency alerts lists.
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.
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.
Status	Current state of the alert. Possible values include Active or Canceled.
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.

Table 4-41. Risk Alerts Data Grid Options (continued)

Option	Description
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.
Owner	Name of the user who owns the alert.
Created On	Date and time when 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.

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.

How Efficiency Alerts Work

All the efficiency alerts generated for you managed objects appear in the list.

You can manage the alerts in the list using the toolbar options, click the alert name to see the alert details for the affected object, or click the name of the object on which the alert was generated to see the object details.

Efficiency Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to cancel, suspend, or manage ownership. You can select multiple rows in the list using Shift+click, Control +click. Use the data grid to view the alerts. You can click the alert name to view the alert details or object name to view the object details.

Table 4-42. Efficiency Alerts Toolbar Options

Option	Description
Open in external application	Actions you can run on the selected object. For example, Open Virtual Machine in vSphere Client.
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.
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.
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.

The Efficiency Alerts data grid provides a list of generated alerts that you use to resolve problems in your environment.

Table 4-43. Efficiency Alerts Data Grid Options

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**. The possible values include: ■ Critical ■ Immediate ■ Warning ■ Information By default, alerts are sorted by criticality. Presorting the alerts list by criticality displays critical alerts at the top of the list. If you change the sort order, the sort is saved with your preferences in the global alerts list, and the Health, Risk, and Efficiency alerts lists.
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.
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.
Status	Current state of the alert. Possible values include Active or Canceled.
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.

Table 4-43. Efficiency Alerts Data Grid Options (continued)

Option	Description
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.
Owner	Name of the user who owns the alert.
Created On	Date and time when 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.

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 **Alert Settings > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-44. 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. ■ 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.
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.

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 generally 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 **Alert Settings > Symptom Definitions > Metric / Property**. Click the plus sign 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 4-45. 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.
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 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.

Table 4-45. Symptoms Workspace Options for Metrics and Super Metrics (continued)

Option	Description
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 4-45. 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.

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.
	  ■ 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.
	■ 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.

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 **Alert Settings > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-46. 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. ■ 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.

Table 4-46. Property Symptoms Definitions Options (continued)

Option	Description
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.

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 **Alert Settings > 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 4-47. 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 4-47. 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 **Alert Settings > Symptom Definitions**. Select the **Message Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-48. 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. ■ 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.

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 **Alert Settings > 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 4-49. 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.
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 **Alert Settings > Symptom Definitions**. Select the **Fault** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-50. 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. ■ 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.

Table 4-50. Fault Symptoms Definitions Options (continued)

Option	Description
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 resource type.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.

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 **Alert Settings > Symptom Definitions**. Click the **Fault** tab and **Fault 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 4-51. 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 4-51. 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 **Alert Settings > Symptom Definitions**. Click the **Metric Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-52. 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. ■ 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 Metric	Selected event metric based on resource 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.

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 **Alert Settings > 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 4-53. 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.
Metric Event list	List of the metric events for the selected base object type.

Table 4-53. Symptoms Workspace Options for Metric Events (continued)

Option	Description
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 4-54. 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 **Alert Settings > Recommendations**.

You can also define recommendations when you create an alert definition.

Table 4-55. Recommendations Overview Options

Option	Description
Toolbar options	Use the toolbar options to manage your recommendations. ■ Add. Add a recommendation. ■ 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.

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 select **Recommendations** from the **Alert Settings** drop-down menu in the left pane. Click **Add** to create a recommendation.

You can also define recommendations when you define alerts.

Table 4-56. 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.
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.

The actions named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear, but 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 actions named `Set Memory for VM Power Off Allowed`, `Set CPU Count for VM Power Off Allowed`, and `Set CPU Count and Memory for VM Power Off Allowed` are also not visible except in the alert recommendations, and 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 **Alert Settings > Alert Definitions**.

Table 4-57. Alert Definition Options

Option	Description
Toolbar options	Use the toolbar options to manage your alert definitions. ■ Add. Add an alert definition. ■ 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 or 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.
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.

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.

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 **Alert Settings > Alert Definitions**. Click the plus sign to add a definition, or click the pencil 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, and the recommendations that might resolve the alert.

- Alert Definition Workspace Name and Description

  The name and description of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager.

- Alert Definition Workspace Base Object Type

  The base object type is the object type on which the alert is generated in vRealize Operations Manager when a symptom condition is found to be true.

- Alert Definition Workspace Alert Impact

  The alert impact specifies the urgency of the alert, determines which badge the alert affects, how critical the alert is to the functioning of your environment, and how it is classified when you or the system processes a generated alert.

- 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 Name and Description

The name and description of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager.

Where You Define Name and Description

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click**Alert Settings > Alert Definitions**. Click the plus sign to add a definition, or click the pencil to edit the selected definition. In the workspace, on the left, click **Name and Description**.

Table 4-58. Alert Definition Name and Description Options

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.

Alert Definition Workspace Base Object Type

The base object type is the object type on which the alert is generated in vRealize Operations Manager when a symptom condition is found to be true.

Where You Define the Base Object Type

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Alert Settings > Alert Definitions**. Click the plus sign to add a definition, or click the pencil to edit the selected definition. In the workspace, on the left, click **Base Object Type**.

Alert Details

Notes

Table 4-59. Base Object Type Options

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.

Alert Definition Workspace Alert Impact

The alert impact specifies the urgency of the alert, determines which badge the alert affects, how critical the alert is to the functioning of your environment, and how it is classified when you or the system processes a generated alert.

Where You Define the Alert Impact

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click**Alert Settings > Alert Definitions**. Click the plus sign to add a definition, or click the pencil to edit the selected definition. In the workspace, on the left, click **Alert Impact**.

Table 4-60. Alert Impact Options

Option	Description
Impact	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 4-60. Alert Impact Options (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.

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.

Where You Define the Symptom Definitions

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Alert Settings > Alert Definitions**. Click the plus sign to add a definition, or click the pencil to edit the selected definition. In the workspace, on the left, click **Add Symptom Definitions**.

Add Symptoms Definitions Options

To add symptom definitions, you use the left pane to select your symptoms. You use the workspace on the right to define the point at which the symptoms or symptom sets are true. You also use the workspace to specify whether all or any of the symptoms or symptom sets must be true to generate an alert.

Table 4-61. Add Symptoms Selection Options

Option	Description
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 4-61. Add Symptoms Selection Options (continued)

Option	Description
Symptom Definition Type	Select the type of symptom definition that you are adding for the current Defined On object type. ■ Metric / Supermetric. Add symptoms that use metric and super metric symptoms. These metrics are based on the operational or performance values that vRealize Operations Managercollects from target objects in your environment. ■ Property. Add symptoms that use property symptoms. These symptoms are based on the configuration properties that vRealize Operations Manager collects from the 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.
Add symptom button	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 defined 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 and the red **x** that appears next to the filter name.

Table 4-61. Add Symptoms Selection Options (continued)

Option	Description
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 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 4-62. Symptom Sets in the Alert Definition Workspace

Option	Description
Alert Definition Summary	The currently configured information for the alert definition. Use the information as reference when you create alert definitions.
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.

Table 4-62. Symptom Sets in the Alert Definition Workspace (continued)

Option	Description
Match {operator} of the following symptom sets	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.
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 **Negate This Symptom Condition** icon to the left of the symptom name. Although you can configure symptom criticality, if you negate a symptom, it does not have an associated criticality that affects the criticality of generated alerts.

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.

Where You Find the Add Recommendation Options

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click**Alert Settings > Alert Definitions**. Click the plus sign to add a definition or click the pencil to edit the selected definition..In the workspace, on the left, click **Add Recommendations**.

Table 4-63. Add Recommendations Options in the Alert Definition Workspace

Option	Description
Add recommendation	If recommendations that you need to resolve the symptoms in the problem do not exist, you can create them.
Quick filter (Name)	Limits the list based on the text you type.
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 in the table.

Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for vRealize Operations Manager to alert you. When the alert is triggered on your host system, vRealize Operations Manager alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have vRealize Operations Manager notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

Procedure

1. In the menu, click **Alerts** and then in the left pane, select **Alert Settings > Alert Definitions**.
2. Enter **`capacity`** in the search text box.

   Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.

3 Click the plus sign to create a new capacity alert definition for your host systems.

a In the alert definition workspace, for the Name and Description, enter **`Hosts – Alert on Capacity Exceeded`**.

b For the Base Object Type, select **vCenter Adapter > Host System**

c For the Alert Impact, select the following options.

Option	Selection
Impact	Select **Risk**.
Criticality	Select **Immediate**.
Alert Type and Subtype	Select **Application : Capacity**.
Wait Cycle	Select **1**.
Cancel Cycle	Select **1**.

d For Add Symptom Definitions, select the following options.

Option	Selection
Defined On	Select **Self**.
Symptom Definition Type	Select **Metric / Supermetric**.
Quick filter (Name)	Enter **`capacity`**.

e From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the right pane.

In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.

f For Add Recommendations, enter **`virtual machine`** in the quick filter text box.

g Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the right pane.

This recommendation is set to Priority 1.

4 Click **Save** to save the alert definition.

Your new alert appears in the list of alert definitions.

Results

You have added an alert definition to have vRealize Operations Manager alert you when the capacity of your host systems begins to run out.

Alert Definition Best Practices

As you create alert definitions for your environment, apply consistent best practices so that you optimize alert behavior for your monitored objects.

Alert Definitions Naming and Description

The alert definition name is the short name that appears in the following places:

- In data grids when alerts are generated
- In outbound alert notifications, including the email notifications that are sent when outbound alerts and notifications are configured in your environment

Ensure that you provide an informative name that clearly states the reported problem. Your users can evaluate alerts based on the alert definition name.

The alert definition description is the text that appears in the alert definition details and the outbound alerts. Ensure that you provide a useful description that helps your users understand the problem that generated the alert.

Wait and Cancel Cycle

The wait cycle setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition goes into effect after the wait cycle for the symptom definition results in a triggered symptom. In most alert definitions you configure the sensitivity at the symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that the alert is immediately generated after all of the symptoms are triggered at the desired symptom sensitivity level.

The cancel cycle setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition goes into 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 Manageron the target systems.

Creating and Managing vRealize Operations Manager Alert Notifications

When alerts are generated in vRealize Operations Manager, they appear in the alert details and object details, but you can also configure vRealize Operations Manager to send your alerts to outside applications using one or more outbound alert options.

You configure notification options to specify which alerts are sent out for the Standard Email, REST, SNMP, and Log File outbound alert plug-ins. For the other plug-in types, all the alerts are sent when the target outbound alert plug-in is enabled.

The most common outbound alert plug-in is the Standard Email plug-in. You configure the Standard Email plug-in to send notifications to one or more users when an alert is generated that meets the criteria you specify in the notification settings.

List of Outbound Plug-Ins in vRealize Operations Manager

vRealize Operations Manager provides outbound plug-ins. This list includes the name of the plug-in and whether you can filter the outbound data based on your notification settings.

If the plug-in supports configuring notification rules, then you can filter the messages before they are sent to the target system. If the plug-in does not support notifications, all messages are sent to the target system, and you can process them in that application.

If you installed other solutions that include other plug-in options, they appear as a plug-in option with the other plug-ins.

Messages and alerts are sent only when the plug-in is enabled.

Table 4-64. Notification Support for Outbound Plug-Ins

Outbound Plug-In	Configure Notification Rules
Automated Action Plug-in	No The Automated Action plug-in is enabled by default. If automated actions stop working, check the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.
Log File Plug-In	Yes To filter the log file alerts, you can either configure the file named `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

Add Outbound Notification Plug-Ins in vRealize Operations Manager

You add outbound plug-in instances so that you can notify users about alerts or capture alert data outside of vRealize Operations Manager.

You can configure one or more instances of the same plug-in type if you need to direct alert information to multiple target systems.

The Automated Action plug-in is enabled by default. If automated actions stop working, check the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.

- Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts

  You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email vRealize Operations Manager alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals.
- Add a REST Plug-In for vRealize Operations Manager Outbound Alerts

  You add a REST Plug-In so that you can send vRealize Operations Manager alerts to another REST-enabled application where you built a REST Web service to accept these messages.
- Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts

  You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.
- Add a Network Share Plug-In for vRealize Operations Manager Reports

  You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location.

- Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts

  You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

- Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts

  You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts

You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email vRealize Operations Manager alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals.

Prerequisites

Ensure that you have an email user account that you can use as the connection account for the alert notifications. If you choose to require authentication, you must also know the password for this account.

Procedure

1. In the menu, click **Administration** and then in the left pane, click **Management**.
2. Click **Outbound Settings** and click the plus sign to add a plug-in.
3. From the **Plug-In Type** drop-down menu, select **Standard Email Plugin**.

   The dialog box expands to include your SMTP settings.
4. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.
5. Configure the SMTP options appropriate for your environment.

Option	Description
Use Secure Connection	Enables secure communication encryption using SSL/TLS. If you select this option, you must select a method in the **Secure Connection Type** drop-down menu.
Requires Authentication	Enables authentication on the email user account that you use to configure this SMTP instance. If you select this option, you must provide a password for the user account.
SMTP Host	URL or IP address of your email host server.
SMTP Port	Default port SMTP uses to connect with the server.
Secure Connection Type	Select either SSL/TLS as the communication encryption method used in your environment from the drop-down menu. You must select a connection type if you select Use Secure Connection.
User Name	Email user account that is used to connect to the email server.

Option	Description
Password	Password for the connection user account. A password is required if you select Requires Authentication.
Sender Email Address	Email address that appears on the notification message
Sender Name	Displayed name for the sender email address.

6 Click **Save**.

7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the standard email plug-in for outbound SMTP alerts is configured and running.

What to do next

Create notification rules that use the standard email plug-in to send a message to your users about alerts requiring their attention. See User Scenario: Create a vRealize Operations Manager Email Alert Notification .

Add a REST Plug-In for vRealize Operations Manager Outbound Alerts

You add a REST Plug-In so that you can send vRealize Operations Manager alerts to another REST-enabled application where you built a REST Web service to accept these messages.

The REST Plug-In supports enabling an integration, it does not provide an integration. Depending on your target application, you might need an intermediary REST service or some other mechanism that will correlate the alert and object identifiers included in the REST alert output with the identifiers in your target application.

Determine which content type you are delivering to your target application. If you select application/json, the body of the POST or PUT calls that are sent have the following format. Sample data is included.

```
{
   "startDate":1369757346267,
   "criticality":"ALERT_CRITICALITY_LEVEL_WARNING",
   "Risk":4.0,
   "resourceId":"sample-object-uuid",
   "alertId":"sample-alert-uuid",
   "status":"ACTIVE",
   "subType":"ALERT_SUBTYPE_AVAILABILITY_PROBLEM",
   "cancelDate":1369757346267,
   "resourceKind":"sample-object-type",
   "alertName":"Invalid IP Address for connected Leaf Switch",
   "attributeKeyID":5325,
   "Efficiency":1.0,
   "adapterKind":"sample-adapter-type",
   "Health":1.0,
   "type":"ALERT_TYPE_APPLICATION_PROBLEM",
```

```
    "resourceName":"sample-object-name",
    "updateDate":1369757346267,
    "info":"sample-info"
}
```

If you select application/xml, the body of the POST or PUT calls that are sent have the following format:

```
<alert>
  <startDate>1369757346267</startDate>
  <criticality>ALERT_CRITICALITY_LEVEL_WARNING</criticality>
  <Risk>4.0</Risk>
  <resourceId>sample-object-uuid</resourceId>
  <alertId>sample-alert-uuid</alertId>
  <status>ACTIVE</status>
  <subType>ALERT_SUBTYPE_AVAILABILITY_PROBLEM</subType>
  <cancelDate>1369757346267</cancelDate>
  <resourceKind>sample-object-type</resourceKind>
  <alertName>Invalid IP Address for connected Leaf Switch</alertName>
  <attributeKeyId>5325</attributeKeyId>
  <Efficiency>1.0</Efficiency>
  <adapterKind>sample-adapter-type</adapterKind>
  <Health>1.0</Health>
  <type>ALERT_TYPE_APPLICATION_PROBLEM</type>
  <resourceName>sample-object-name</resourceName>
  <updateDate>1369757346267</updateDate>
  <info>sample-info</info>
</alert>
```

Note If the alert is triggered by a non-metric violation, the `attributeKeyID` is omitted from the REST output and is not sent.

If the request is processed as POST, for either JSON or XML, the Web service returns an HTTP status code of 201, which indicates the alert was successfully created at the target. If the request is processed as PUT, the HTTP status code of 202, which indicates the alert was successfully accepted at the target.

Prerequisites

Ensure that you know how and where the alerts sent using the REST plug-in are consumed and processed in your environment, and that you have the appropriate connection information available.

Procedure

1. In the left pane of vRealize Operations Manager, click the **Administration** icon.
2. Click **Outbound Settings** and click the plus sign to add a plug-in.
3. From the **Plug-In Type** drop-down menu, select **Rest Notification Plugin**.

   The dialog box expands to include your REST settings.

4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

5 Configure the Rest options appropriate for your environment.

Option	Description
URL	URL to which you are sending the alerts. The URL must support HTTPS. When an alert is sent to the REST Web server, the plug-in appends `/{alertID}` to the POST or PUT call.
User Name	User account on the target REST system.
Password	User account password.
Content Type	Specify the format for the alert output. ■ 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.
Connection count	Limits the number of simultaneous alerts that are sent to the target REST server. Use this number to ensure that your REST server is not overwhelmed with requests.

6 Click **Save**.

7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the REST plug-in for outbound alerts is configured and running.

What to do next

Create notification rules that use the REST plug-in to send alerts to a REST-enabled application or service in your environment. See User Scenario: Create a vRealize Operations Manager REST Alert Notification.

Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts

You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.

All alerts are added to the log file. You can use other applications to filter and manage the logs.

Prerequisites

Ensure that you have write access to the file system path on the target vRealize Operations Manager nodes.

Procedure

1. In the menu, click **Administration** and then in the left pane, click **Management**.
2. Click **Outbound Settings** and click the plus sign to add a plug-in.
3. From the **Plug-In Type** drop-down menu, select **Log File**.

   The dialog box expands to include your log file settings.
4. In the **Alert Output Folder** text box, enter the folder name.

   If the folder does not exist in the target location, the plug-in creates the folder in the target location. The default target location is: `/usr/lib/vmware-vcops/common/bin/`.
5. Click **Save**.
6. To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the log file plug-in is configured and running.

What to do next

When the plug-in is started, the alerts are logged in the file. Verify that the log files are created in the target directory as the alerts are generated, updated, or canceled.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

1. In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.
2. From the toolbar, click the **Add** icon.
3. From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

   The dialog box expands to include your plug-in instance settings.
4. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.

5 Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.
Password	The password for the domain user account.
Network share root	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.

Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts

You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

You can provide filtering when you define a Notification using an SNMP Trap destination.

Prerequisites

Ensure that you have an SNMP Trap server configured in your environment, and that you know the IP address or host name, port number, and community that it uses.

Procedure

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

2 Click **Outbound Settings** and click the plus sign to add a plug-in.

3 From the **Plug-In Type** drop-down menu, select **SNMP Trap**.

The dialog box expands to include your SNMP trap settings.

4 Type an **Instance Name**.

5 Configure the SNMP trap settings appropriate to your environment.

Option	Description
Destination Host	IP address or fully qualified domain name of the SNMP management system to which you are sending alerts.
Port	Port used to connect to the SNMP management system. Default port is 162.
Community	Text string that allows access to the statistics. SNMP Community strings are used only by devices that support SNMPv3 protocol.
Username	Username to configure SNMP trap settings in your environment. If the username is specified, SNMPv3 is considered as the protocol by the plugin. If left blank, SNMPv2c is considered as the protocol by the plugin.
Authentication Protocol	Authentication algorithms available are SHA-224, SHA-256, SHA-384, SHA-512.
Authentication Password	Authentication password.
Privacy Protocol	Privacy algorithms available are AES192, AES2564.
Privacy Password	Privacy password.

6 Click **Save**.

Results

This instance of the SNMP Trap plug-in is configured and running.

What to do next

When the plug-in is added, Configuring Notifications for receiving the SNMP traps.

Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts

You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

This outbound alert option is useful when you manage the same objects in Server Assurance Manager and in vRealize Operations Manager, and you added the EMC Smarts management pack and configured the solution in vRealize Operations Manager. Although you cannot filter the alerts sent to Service Assurance Manager in vRealize Operations Manager, you can configure the Smarts plug-in to send the alerts to the Smarts Open Integration server. You then configure the Open Integration server to filter the alerts from vRealize Operations Manager, and send only those that pass the filter test to the Smarts Service Assurance Manager service.

Prerequisites

- Verify that you configured the EMC Smarts solution. For documentation regarding EMC Smarts integration, see https://solutionexchange.vmware.com/store.

- Ensure that you have the EMC Smarts Broker and Server Assurance Manager instance host name or IP address, user name, and password.

Procedure

1. In the menu, click **Administration** and then in the left pane, click **Management**.
2. Click **Outbound Settings** and click the plus sign to add a plug-in.
3. From the **Plug-In Type** drop-down menu, select **Smarts SAM Notification**.

   The dialog box expands to include your Smarts settings.
4. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.
5. Configure the Smarts SAM notification settings appropriate for your environment.

Option	Description
Broker	Type the host name or IP address of the EMC Smarts Broker that manages registry for the Server Assurance Manager instance to which you want the notifications sent.
Broker Username	If the Smarts broker is configured as Secure Broker, type the user name for the Broker account.
Broker Password	If the Smarts broker is configured as Secure Broker, type the password for the Broker user account.
SAM Server	Type the host name or IP address of the Server Assurance Manager server to which you are sending the notifications.
User Name	Type the user name for the Server Assurance Manager server instance. This account must have read and write permissions for the notifications on the Smarts server as specified in the SAM Server.
Password	Type the password for the Server Assurance Manager server account.

6. Click **Save**.
7. Modify the Smarts SAM plug-in properties file.
   a. Open the properties file at: `/usr/lib/vmware-vcops/user/plugins/outbound/vcops-smartsalert-plugin/conf/plugin.properties`
   b. Add the following string to the properties file: `# sendByType=APPLICATION::AVAILABILITY,APPLICATION::PERFORMANCE,APPLICATION::CAPACITY,APPLICATION::COMPLIANCE,VIRTUALIZATION::AVAILABILITY,VIRTUALIZATION::PERFORMANCE,VIRTUALIZATION::CAPACITY,VIRTUALIZATION::COMPLIANCE,HARDWARE::AVAILABILITY,HARDWARE::PERFORMANCE,HARDWARE::CAPACITY,HARDWARE::COMPLIANCE,STORAGE::AVAILABILITY,STORAGE::PERFORMANCE,STORAGE::CAPACITY,STORAGE::COMPLIANCE,NETWORK::AVAILABILITY,NETWORK::PERFORMANCE,NETWORK::CAPACITY,NETWORK::COMPLIANCE`
   c. Save the properties file.

8 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the Smarts SAM Notifications plug-in is configured and running.

What to do next

In Smarts Service Assurance Manager, configure your Notification Log Console to filter the alerts from vRealize Operations Manager. To configure the filtering for Service Assurance Manager, see the EMC Smarts Service Assurance Manager documentation.

Filtering Log File Outbound Messages With the TextFilter.xml File

The log file outbound plug-in in vRealize Operations Manager captures alert data. To filter the log file data, you can update the `TextFilter.xml` file to capture only the alerts meeting the filter criteria.

As a vRealize Operations Manager administrator, you want to filter the outbound alert log files based on the alert type and the subtype.

The filters are configured in the `TextFile.xml` file. The file is in the following location:

- vApp or Linux. `/usr/lib/vmware-vcops/user/plugins/outbound/vcops-textfile-plugin/conf`

In the file, use the following format for the filter rule.

```
<FilterRule name="AlertType">
  <AlertTypes>
    <AlertType key="AlertType1:AlertSubType1 " />
    <AlertType key="AlertType2:AlertSubType2 " />
  </AlertTypes>
</FilterRule>
```

For example, the rule to filter based on the Application type and Availability subtype uses this format.

```
<FilterRule name="AlertType">
  <AlertTypes>
    <AlertType key="ALERT_TYPE_APPLICATION_PROBLEM:ALERT_SUBTYPE_AVAILABILITY_PROBLEM " />
  </AlertTypes>
</FilterRule>
```

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 **Outbound Settings**.

Table 4-65. Outbound Settings Options

Option	Description
Toolbar options	Use the toolbar options to manage your Outbound Plug-Ins. ■ Add or Edit. Opens the Outbound Plug-In dialog box where you configure the connection options for the instance. ■ 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 left pane, and click **Outbound Settings**. On the toolbar, click the plus sign to add a plug-in instance, or select a plug-in from the list and click the pencil 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.

Configuring Notifications

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager. You configure notification rules for the supported outbound alerts so that you can filter the alerts that are sent to the selected external system.

You use the notifications list to manage your rules. You then use the notification rules to limit the alerts that are sent to the external system. To use notifications, the supported outbound alert plug-ins must be added and running.

With notification rules, you can limit the data that is sent to the following external systems.

- Standard Email. You can create multiple notification rules for various email recipients based on one or more of the filter selections. If you add recipients but do not add filter selections, all the generated alerts are sent to the recipients.
- REST. You can create a rule to limit alerts that are sent to the target REST system so that you do not need to implement filtering on that target system.
- SNMP Trap. You can configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.
- Log File. You can configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes.

User Scenario: Create a vRealize Operations Manager Email Alert Notification

As a virtual infrastructure administrator, you need vRealize Operations Manager to send email notifications to your advanced network engineers when critical alerts are generated for mmbhost object, the host for many virtual machines that run transactional applications, where no one has yet taken ownership of the alert.

Prerequisites

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Ensure that at least one instance of the standard email plug-in is configured and running. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.

Procedure

1. In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.
2. Click **Notification Settings** and click the plus sign to add a notification rule.
3. In the **Name** text box type a name similar to **`Unclaimed Critical Alerts for mmbhost`**.
4. In the Method area, select **Standard Email Plug-In** from the drop-down menu, and select the configured instance of the email plug-in.

5 Configure the email options.

   a In the **Recipients** text box, type the email addresses of the members of your advance engineering team, separating the addresses with a semi-colon (;).

   b To send a second notification if the alert is still active after a specified amount of time, type the number of minutes in the **Notify again** text box.

   c Type number of notifications that are sent to users in the **Max Notifications** text box.

6 Configure the scope of filtering criteria.

   a From the **Scope** drop-down menu, select **Object**.

   b Click **Click to select Object** and type the name of the object.

   In this example, type **mmbhost**.

   c Locate and select the object in the list, and click **Select**.

7 Configure the Notification Trigger.

   a From the **Notification Trigger** drop-down menu, select **Impact**.

   b From the adjacent drop-down menu, select **Health**.

8 In the Criticality area, click **Critical**.

9 Expand the Advanced Filters and from the **Alert States** drop-down menu, select **Open**.

   The Open state indicates that no engineer or administrator has taken ownership of the alert.

10 Click **Save**.

Results

You created a notification rule that sends an email message to the members of your advance network engineering team when any critical alerts are generated for the mmbhost object and the alert is not claimed by an engineer. This email reminds them to look at the alert, take ownership of it, and work to resolve the triggering symptoms.

What to do next

Respond to alert email notifications. See User Scenario: An Alert Arrives in Your Inbox.

User Scenario: Create a vRealize Operations Manager REST Alert Notification

As a virtual infrastructure administrator, you need vRealize Operations Manager to send alerts in JSON or XML to a REST-enabled application that has REST Web service that accepts these messages. You want only alerts where the virtualization alerts that affect availability alert types go to this outside application. You can then use the provided information to initiate a remediation process in that application to address the problem indicated by the alert.

The notification configuration limits the alerts sent to the outbound alert instance to those matching the notification criteria.

Prerequisites

- Verify that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Verify that at least one instance of the REST plug-in is configured and running. See Add a REST Plug-In for vRealize Operations Manager Outbound Alerts.

Procedure

1. In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.
2. Click **Notifications** and click the plus sign to add a notification rule.
3. In the **Name** text box type a name similar to **`Virtualization Alerts for Availability`**.
4. In the Method area, select **REST Plug-In** from the drop-down menu, and select the configured instance of the email plug-in.
5. Configure the Notification Trigger.
   a. From the **Notification Trigger** drop-down menu, select **Alert Type**.
   b. Click **Click to select Alert type/subtype** and select **Virtualization/Hypervisor Alerts Availability**.
6. In the Criticality area, click **Warning**.
7. Expand the Advanced Filters and from the **Alert Status** drop-down menu, select **New**.

   The New status indicates that the alert is new to the system and not updated.
8. Click **Save**.

Results

You created a notification rule that sends the alert text to the target REST-enabled system. Only the alerts where the configured alert impact is Virtualization/Hypervisor Availability and where the alert is configured as a warning are sent to the target instance using the REST plug-in.

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, select **Alerts** in the menu, and click **Notifications Settings** from the left pane.

Table 4-66. Notifications Options

Option	Description
Toolbar options	Use the toolbar options to manage your notification rules. ■ Add or Edit. Opens the Rule dialog box where you configure the filtering options for the notification rule. ■ Delete. Removes the selected rule.
Rule Name	Name you assigned when you created the notification 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.
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.

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 Manager 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 **Alert Settings > Notification Settings**. On the toolbar, click the plus sign to add a rule, or select a rule and click the pencil to edit the exiting rule.

Table 4-67. 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 plug-in. Select one of the configured outbound alert plug-in types: Standard Email, REST, SNMP Trap, and Log File. ■ Instance. Select the configured instance for the type of plug-in. ■ Recipients. (Standard Email Plug-In only) 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. (Standard Email Plug-In only) Number of minutes between notifications messages for active alerts. Leave the text box empty to send only one message per alert. ■ Max Notifications. (Standard Email Plug-In only) 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. (Standard Email Plug-In only) 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.
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.
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.
Alert States	Managed state of the alert, either opened, assigned, or suspended.
Alert Status	Current state of the alert, either canceled, updated, or new.
Collectors	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.

Create an Alert Definition for Department Objects

As a virtual infrastructure administrator, you are responsible for the virtual machines and hosts that the accounting department uses. You can create alerts to manage the accounting department objects.

You received several complaints from your users about delays when they are using their accounting applications. Using vRealize Operations Manager, you identified the problem as related to CPU allocations and workloads. To better manage the problem, you create an alert definition with tighter symptom parameters so that you can track the alerts and identify problems before your users encounter further problems.

Using this scenario, you create a monitoring system that monitors your accounting objects and provides timely notifications when problems occur.

Add Description and Base Object to Alert Definition

To create an alert to monitor the CPUs for the accounting department virtual machines and monitor host memory for the hosts on which they operate, you begin by describing the alert.

When you name the alert definition and define alert impact information, you specify how the information about the alert appears in vRealize Operations Manager. The base object is the object around which the alert definition is created. The symptoms can be for the base object and for related objects.

Procedure

1. In the menu, click **Alerts** and then in the left pane, click **Alert Settings > Alert Definitions**.
2. Click the plus sign to add a definition.
3. Type a name and description.

   In this scenario, type **`Acct VM CPU early warning`** as the alert name, which is a quick overview of the problem. The description, which is a detailed overview. should provide information that is as useful as possible. When the alert is generated, this name and description appears in the alert list and in the notification.
4. Click **Base Object Type**.
5. From the drop-down menu, expand **vCenter Adapter** and select **Host System**.

   This alert is based on host systems because you want an alert that acts as an early warning to possible CPU stress on the virtual machines used in the accounting department. By using host systems as the based object type, you can respond to the alert symptom for the virtual machines with bulk actions rather than responding to an alert for each virtual machine.
6. Click **Alert Impact** and configure the metadata for this alert definition.
   a. From the **Impact** drop-down menu, select **Risk**.

      This alert indicates a potential problem and requires attention in the near future.

   b. From the **Criticality** drop-down menu, select **Immediate**.

      As a Risk alert, which is indicative of a future problem, you still want to give it a high criticality so that it is ranked for correct processing. Because it is designed as an early warning, this configuration provides a built-in buffer that makes it an immediate risk rather than a critical risk.

c From the **Alert Type and Subtype** drop-down menu, expand **Virtualization/Hypervisor** and select **Performance**.

d To ensure that the alert is generated during the first collection cycle after the symptoms become true, set the **Wait Cycle** to **1**.

e To ensure that the an alert is removed as soon as the symptoms are no longer triggered, set the **Cancel Cycle** to **1**.

The alert is canceled in the next collection cycle if the symptoms are no long true.

These alert impact options help you identify and prioritize alerts as they are generated.

Results

You started an alert definition where you provided the name and description, selected host system as the base object type, and defined the data that appears when the alert generated.

What to do next

Continue in the workspace, adding symptoms to your alert definition. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

Add a Virtual Machine CPU Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add symptoms to your vRealize Operations Manager alert definition after you provide the basic descriptive information for the alert. The first symptom you add is related to CPU usage on virtual machines. You later use a policy and group to apply alert to the accounting virtual machines.

This scenario has two symptoms, one for the accounting virtual machines and one to monitor the hosts on which the virtual machines operate.

Prerequisites

Begin configuring the alert definition. See Add Description and Base Object to Alert Definition.

Procedure

1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Add Symptom Definitions** and configure the symptoms.

2 Begin configuring the symptom set related to virtual machines CPU usage.

a From the **Defined On** drop-down menu, select **Child**.

b From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.

c From the **Symptom Definition Type** drop-down menu, select **Metric / Supermetric**.

d Click the **Add** button to open the **Add Symptom Definition** workspace window.

3 Configure the virtual machine CPU usage symptom in the **Add Symptom Definition** workspace window.

   a From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Virtual Machine**.

      The collected metrics for virtual machines appears in the list.

   b In the metrics list **Search** text box, which searches the metric names, type **usage**.

   c In the list, expand **CPU** and drag **Usage (%)** to the workspace on the right.

   d From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.

   e In the **Symptom Definition Name** text box, type a name similar to **VM CPU Usage above trend**.

   f From the criticality drop-down menu, select **Warning**.

   g From the threshold drop-down menu, select **Above Threshold**.

   h Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for 3 collection cycles before the symptom is triggered. This wait avoids triggering the symptom when there is a short spike in CPU usage.

   i Click **Save**.

   The dynamic symptom, which identifies when the usage is above the tracked trend, is added to the symptom list.

4 In the **Alert Definition Workspace** window, drag **VM CPU Usage above trend** from the symptom definition list to the symptom workspace on the right.

   The Child-Virtual Machine symptom set is added to the symptom workspace.

5 In the symptoms set, configure the triggering condition so that when the symptom is true on half of the virtual machines in the group to which this alert definition is applied, the symptom set is true.

   a From the value operator drop-down menu, select **>**.

   b In the value text box, enter **50**.

   c From the value type drop-down menu, select **Percent**.

Results

You defined the first symptom set for the alert definition.

What to do next

Add the host memory usage symptom to the alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

Add a Host Memory Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add a second symptom to your vRealize Operations Manager alert definition after you add the first symptom. The second symptom is related to host memory usage for the hosts on which the accounting virtual machines operate.

Prerequisites

Add the virtual machine CPU usage symptom. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

Procedure

1. In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Add Symptom Definitions**.
2. Configure the symptom related to host systems for the virtual machines.
   a. From the **Defined On** drop-down menu, select **Self**.
   b. From the **Symptom Definition Type** drop-down menu, select **Metric / Supermetric**.
   c. Click the **Add** button to configure the new symptom.
3. Configure the host system symptom in the **Add Symptom Definition** workspace window.
   a. From the **Base Object Type** drop-down menu, expand **vCenter Adapters** and select **Host System**.
   b. In the metrics list, expand **Memory** and drag **Usage (%)** to the workspace on the right.
   c. From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.
   d. In the **Symptom Definition Name** text box, enter a name similar to **`Host memory usage above trend`**.
   e. From the criticality drop-down menu, select **Warning**.
   f. From the threshold drop-down menu, select **Above Threshold**.
   g. Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for three collection cycles before the symptom is triggered. This wait avoids triggering the symptom when a short spike occurs in host memory usage.
   h. Click **Save**.

The dynamic symptom identifies when the hosts on which the accounting virtual machines run are operating above the tracked trend for memory usage.

The dynamic symptom is added to the symptom list.

4 In the **Alert Definition Workspace** window, drag **Host memory usage above trend** from the symptoms list to the symptom workspace on the right.

The Self-Host System symptom set is added to the symptom workspace.

5 On the Self-Host System symptom set, from the value type drop-down menu for **This Symptom set is true when**, select **Any**.

With this configuration, when any of the hosts running accounting virtual machines exhibit memory usage that is above the analyzed trend, the symptom condition is true.

6 At the top of the symptom set list, from the **Match {operator} of the following symptoms** drop-down menu, select **Any**.

With this configuration, if either of the two symptom sets, virtual machine CPU usage or the host memory, are triggered, an alert is generated for the host.

Results

You defined the second symptom set for the alert definition and configured how the two symptom sets are evaluated to determine when the alert is generated.

What to do next

Add recommendations to your alert definition so that you and your engineers know how to resolve the alert when it is generated. See Add Recommendations to the Alert Definition.

Add Recommendations to the Alert Definition

To resolve a generated alert for the accounting department's virtual machines, you provide recommendations so that you or other engineers have the information you need to resolve the alert before your users encounter performance problems.

As part of the alert definition, you add recommendations that include actions that you run from vRealize Operations Manager and instructions for making changes in vCenter Server that resolve the generated alert.

Prerequisites

Add symptoms to your alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

Procedure

1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, **Alert Impact**, and **Add Symptom Definitions**, click **Add Recommendations** and add the recommended actions and instructions.

2 Click **Add** and select an action recommendation to resolve the virtual machine alerts.

   a In the **New Recommendation** text box, enter a description of the action similar to `Add CPUs to virtual machines`.

   b From the **Actions** drop-down menu, select **Set CPU Count for VM**.

   c Click **Save**.

3 Click **Add** and provide an instructive recommendation to resolve host memory problems similar to this example.

   `If this host is part of a DRS cluster, check the DRS settings to verify that the load balancing setting are configured correctly. If necessary, manually vMotion the virtual machines.`

4 Click **Add** and provide an instructive recommendation to resolve host memory alerts.

   a Enter a description of the recommendation similar to this example.

   `If this is a standalone host, add more memory to the host.`

   b To make the URL a hyperlink in the instructions, copy the URL, for example, https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html, to your clipboard.

   c Highlight the text in the text box and click **Create a hyperlink**.

   d Paste the URL in the **Create a hyperlink** text box and click **OK**.

   e Click **Save**.

5 In the **Alert Definition Workspace**, drag **Add CPUs to virtual machines**, **If this host is part of a DRS cluster**, and the **If this is a standalone host** recommendations from the list to the recommendation workspace in the order presented.

6 Click **Save**.

Results

You provided the recommended actions and instructions to resolve the alert when it is generated. One of the recommendations resolves the virtual machine CPU usage problem and the other resolves the host memory problem.

What to do next

Create a group of objects to use to manage your accounting objects. See Create a Custom Accounting Department Group.

Create a Custom Accounting Department Group

To manage, monitor, and apply policies to the accounting objects as a group, you create a custom object group.

Prerequisites

Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.

Procedure

1. In the menu, click **Environment** and click the **Custom Groups** tab.
2. Click on the plus sign 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.

3 Click **Add New Policy**.

4 Type a name similar to **Accounting Objects Alerts Policy** and provide a useful description similar to the following example.

```
This policy is configured to generate alerts when
Accounting VMs and Hosts group objects are above trended
 CPU or memory usage.
```

5 Click **Select Base Policies** and select **Default Policy** from the **Start with** drop-down menu.

6 On the left, click **Customize Alert / Symptom Definitions** and disable all the alert definitions except the new Acct VM CPU early warning alert.

   a In the Alert Definitions area, click **Actions** and select **Select All**.

      The alerts on the current page are selected.

   b Click **Actions** and select **Disable**.

      The alerts indicate Disabled in the State column.

   c Repeat the process on each page of the alerts list.

   d Select **Acct VM CPU early warning** in the list, click **Actions** and select **Enable**.

      The Acct VM CPU early warning alert is now enabled.

7 On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.

8 Click **Save**.

Results

You created a policy where the accounting alert definition exists in a custom policy that is applied only to the virtual machines and hosts for the accounting department.

What to do next

Create an email notification so that you learn about alerts even you when you are not actively monitoring vRealize Operations Manager. See Configure Notifications for the Department Alert.

Configure Notifications for the Department Alert

To receive an email notification when the accounting alert is generated, rather than relying on your ability to generally monitor the accounting department objects in vRealize Operations Manager, you create notification rules.

Creating an email notification when accounting alerts are triggered is an optional process, but it provides you with the alert even when you are not currently working in vRealize Operations Manager.

Prerequisites

- Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.

- Verify that standard email outbound alerts are configured in your system. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.

Procedure

1 In the menu, click **Alerts** and then in the left pane, click **Alert Settings**.

2 Click **Notification Settings** and click the plus sign to add a notification rule.

3 Configure the communication options.

   a In the **Name** text box, type a name similar to **`Acct Dept VMs or Hosts Alerts`**.

   b From the **Select Plug-In Type** drop-down menu, select **StandardEmailPlugin**.

   c From the **Select Instance** drop-down menu, select the standard email instance that is configured to send messages.

   d In the **Recipients** text box, type your email address and the addresses of other recipients responsible for the accounting department alerts. Use a semicolon between recipients.

   e Leave the **Notify again** text box blank.

   If you do not provide a value, the email notice is sent only once. This alert is a Risk alert and is intended as an early warning rather than requiring an immediate response.

   You configured the name of the notification when it is sent to you and the method that is used to send the message.

4 In the Filtering Criteria area, configure the accounting alert notification trigger.

   a From the **Notification Trigger** drop-down menu, select **Alert Definition**.

   b Click **Click to select Alert Definition**.

   c Select **Acct VM CPU early warning** and click **Select**.

5 Click **Save**.

Results

You created a notification rule that sends you and your designated engineers an email message when this alert is generated for your accounting department alert definition.

What to do next

Create a dashboard with alert-related widgets so that you can monitor alerts for the accounting object group. See Create a Dashboard to Monitor Department Objects.

Create a Dashboard to Monitor Department Objects

To monitor all the alerts related to the accounting department object group, you create a dashboard that includes the alert list and other widgets. The dashboard provides the alert data in a single location for all related objects.

Creating a dashboard to monitor the accounting virtual machines and related hosts is an optional process, but it provides you with a focused view of the accounting object group alerts and objects.

Prerequisites

Create an object group for the accounting department virtual machines and related objects. See Create a Custom Accounting Department Group.

Procedure

1. In the menu, click **Dashboards > Actions > Create Dashboard**.
2. In the Dashboard Configuration definition area, type a tab name similar to **`Accounting VMs and Hosts`** and configure the layout options.
3. Click **Widget List** and drag the following widgets to the workspace.
   - **Alert List**
   - **Efficiency**
   - **Health**
   - **Risk**
   - **Top Alerts**
   - **Alert Volume**

   The blank widgets are added to the workspace. To change the order in which they appear, you can drag them to a different location in the workspace.
4. On the Alert List widget title bar, click **Edit Widget** and configure the settings.
   a. In the **Title** text box, change the title to **`Acct Dept Alert List`**.
   b. For the **Refresh Content** option, select **On**.
   c. Type **`Accounting`** in the **Search** text box and click **Search**.

      The Accounting value corresponds to the name of the object group for the accounting department virtual machines and related hosts.
   d. In the filtered resource list, select the **Accounting VMs and Hosts** group.

      The Accounting VMs and Hosts group is identified in the Selected Resource text box.
   e. Click **OK**.

   The Acct Dept Alert List is now configured to display alerts for the Accounting VMs and Hosts group objects.

5 Click **Widget Interactions** and configure the following interactions.

   a For Acct Dept Alert List, leave the selected resources blank.

   b For Top Alerts, Health, Risk, Efficiency, and Alert Volume select **Acct Dept Alert List** from the **Selected Resources** drop-down menu.

   c Click **Apply Interactions**.

   With the widget interaction configured in this way, the select alert in the Acct Dept Alert List is the source for the data in the other widgets. When you select an alert in the alert list, the Health, Risk, and Efficiency widgets display alerts for that object, Top Alerts displays the topic issues affecting the health of the object, and Alert Volume displays an alert trend chart.

6 Click **Save**.

Results

You created a dashboard that displays the alerts related to the accounting virtual machines and hosts group, including the Risk alert you created.

Alerts Group

For easy and better management of alerts, you can arrange them as a group as per your requirement.

It is complicated to identify a problem in large environments as you receive different kind of alerts. To manage alerts easily, group them by their definitions.

For example, there are 1000 alerts in your system. To identify different types of alerts, group them based on their alert definitions. It is also easy to detect the alert having the highest severity in the group.

When you group alerts, you can see the number of times the alerts having the same alert definition are triggered. By grouping alerts, you can perform the following tasks easily and quickly:

- Find the noisiest alert: The alert that has triggered maximum number of times is known as the noisiest alert. Once you find it, you can disable it to avoid further noise.
- Filter alerts: You can filter alerts based on a substring in alert definitions. The result shows the group of alerts that contain the substring.

Note

- If you cancel or disable an alert group, the alerts are not canceled instantly. It might take some time if the group is large.
- Only one group can be expanded at a time.
- The number next to the group denotes the number of alerts in that particular group.
- The criticality sign indicates the highest level of severity of an alert in a group.

Grouping Alerts

You can group alerts by time, criticality, definition, and object type.

To group alerts:

Procedure

1. In the menu, click **Alerts**.
2. Select from the various options available from the **Group By** drop-down menu.

Disable Alerts

In an alerts group, you can disable an alert by a single click.

To disable an alert, in the menu, click **Alerts** and then in the left pane, click **All Alerts**. Select the alert name from the data grid, and click **Actions > Disable**.

The alerts can be disabled by two methods:

- Disable Alert in All Policies: You disable the alert for all the objects for all the policies.
- Disable Alert in Selected Policies: You disable the alert for the objects having the selected policy. Note that this method will work only for objects with alerts.

Configuring Actions

Actions are the ability to update objects or read data about objects in monitored systems, and are commonly provided in vRealize Operations Manager as part of a solution. The actions added by solutions are available from the object Actions menu, list and view menus, including some dashboard widgets, and can be added to alert definition recommendations.

The possible actions include read actions and update actions.

The read actions retrieve data from the target objects.

The update actions modifies the target objects. For example, you can configure an alert definition to notify you when a virtual machine is experiencing memory issues. Add an action in the recommendations that runs the Set Memory for Virtual Machine action. This action increases the memory and resolves the likely cause of the alert.

To see or use the actions for your vCenter Server objects, you must enable actions in the vCenter Adapter for each monitored vCenter Server instance. Actions can only be viewed and accessed if you have the required permissions.

List of vRealize Operations Manager Actions

The list of actions includes the name of the action, the objects that each one modifies, and the object levels at which you can run the action. You use this information to ensure that you correctly apply the actions as alert recommendations and when the actions are available in the **Actions** menu.

Actions and Modified Objects

vRealize Operations Manager actions make changes to objects in your managed vCenter Server instances.

When you grant a user access to actions in vRealize Operations Manager, that user can take the granted action on any object that vRealize Operations Manager manages, and not only on objects that the user can access outside of vRealize Operations Manager.

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 4-68. 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

Table 4-68. vRealize Operations Manager Actions Affected Objects (continued)

Action	Modified Object	Object Levels
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
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

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 **Alert Settings > Actions**.

Table 4-69. 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.

Table 4-69. Actions Overview Options (continued)

Option	Description
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.

The actions named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear, but 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 actions named `Set Memory for VM Power Off Allowed`, `Set CPU Count for VM Power Off Allowed`, and `Set CPU Count and Memory for VM Power Off Allowed` are also not visible except in the alert recommendations, and 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 > Alert Settings > Alert Definitions**.
- Create, edit, and import recommendation definitions in **Alerts > Alert Settings > Recommendations**.

Important You set the permissions used to run the actions separately from the alert and recommendation definition. Anyone who can modify alerts, recommendations, and policies can also automate the action, even if they do not have permission to run the action.

For example, if you do not have access to the Power Off VM action, but you can create and modify alerts and recommendations, you can see the Power Off VM action and assign it to an alert recommendation. Then, if you automate the action in your policy, vRealize Operations Manager uses the `automationAdmin` user to run the action.

Example Action Supported for Automation

For the Alert Definition named `Virtual machine has chronic high CPU workload leading to CPU stress`, you can automate the action named `Set CPU Count for VM`.

When CPU stress on your virtual machines exceeds a critical, immediate, or warning level, the alert triggers the recommended action without user intervention.

Integration of Actions with vRealize Automation

vRealize Operations Manager restricts actions on objects that vRealize Automation manages, so that the actions do not violate any constraints set forth by vRealize Automation.

When objects in your environment are managed by vRealize Automation, actions in vRealize Operations Manager are not available on those objects. For example, if a host or parent object is being managed by vRealize Automation, actions are not available on that object.

This behavior is true for all actions, including **Power Off VM**, **Move VM**, **Rebalance Container**, and so on.

You cannot turn on or turn off the exclusion of actions on vRealize Automation managed objects.

Actions Determine Whether Objects Are Managed

Actions check the objects in the vRealize Automation managed resource container to determine which objects are being managed by vRealize Automation.

- Actions such as Rebalance Container check the child objects of the data center container or custom data center container to determine whether the objects are managed by vRealize Automation. If the objects are being managed, the action does not appear on those objects.
- The Move VM action checks whether the virtual machine to be moved is being managed by vRealize Automation.

Is the Virtual Machine Managed?	Result of Move VM Action
Yes	The Move VM action does not appear in the vRealize Operations Manager user interface for that virtual machine.
No	The Move VM action moves the virtual machine to a new host, datastore, or new host and datastore. The Move VM action does not check whether the new host or datastore is being managed by vRealize Automation.

- The Delete Snapshots action checks whether the virtual machine or datastore is being managed by vRealize Automation.

Actions on Objects that vRealize Automation Does Not Manage

For a host or parent object that is not managed by vRealize Automation, only the virtual machines that are not being managed by vRealize Automation appear in the action dialog, and you can only take action on the virtual machines that are not being managed by vRealize Automation. If all child objects are being managed by vRealize Automation, the user interface displays the message `No objects are eligible for the selected action.`

If You Attempt to Run an Action on Multiple Objects

If you select multiple objects and attempt to run an action, such as Power Off VM, only the objects that are not being managed by vRealize Automation, which might include a subset of the virtual machines, appear in the Power Off VM action dialog box.

Working With Actions That Use Power Off Allowed

Some of the actions provided with vRealize Operations Manager require the virtual machines to shut down or power off, depending on the configuration of the target machines, to run the actions. You should understand the impact of the Power Off Allowed option before running the actions so that you select the best options for your target virtual machines.

Power Off and Shut Down

The actions that you can run on your vCenter Server instances include actions that shut down virtual machines and actions that power off virtual machines. It also includes actions where the virtual machine must be in a powered off state to complete the action. Whether the virtual machine 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 the virtual machine off without regard for the state of the guest operating system. In this case, if the virtual machine is running applications, your user could 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 virtual machine, some operating systems support the actions if the Hot Plug is configured on the virtual machine, but 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 are 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 are installed and running, the virtual machine is shut down before completing the action.
- If VMware Tools are 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 increasing 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 4-70. 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 4-71. 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 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 4-72. 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.
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 Policies

To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

Policies

A policy is a set of rules that you define for vRealize Operations Manager to use to analyze and display information about the objects in your environment. You can create, modify, and administer policies to determine how vRealize Operations Manager displays data in dashboards, views, and reports.

How Policies Relate to Your Environment

vRealize Operations Manager policies support the operational decisions established for your IT infrastructure and business units. With policies, you control what data vRealize Operations Manager collects and reports on for specific objects in your environment. Each policy can inherit settings from other policies, and you can customize and override various analysis settings, alert definitions, and symptom definitions for specific object types, to support the service Level agreements and business priorities established for your environment.

When you manage policies, you must understand the operational priorities for your environment, and the tolerances for alerts and symptoms to meet the requirements for your business critical applications. Then, you can configure the policies so that you apply the correct policy and threshold settings for your production and test environments.

Policies define the settings that vRealize Operations Manager applies to your objects when it collects data from your environment. vRealize Operations Manager applies policies to newly discovered objects, such as the objects in an object group. For example, you have an existing VMware adapter instance, and you apply a specific policy to the group named World. When a user adds a new virtual machine to the vCenter Server instance, the VMware adapter reports the virtual machine object to vRealize Operations Manager. The VMware adapter applies the same policy to that object, because it is a member of the World object group.

To implement capacity policy settings, you must understand the requirements and tolerances for your environment, such as CPU use. Then, you can configure your object groups and policies according to your environment.

- For a production environment policy, a good practice is to configure higher performance settings, and to account for peak use times.
- For a test environment policy, a good practice is to configure higher utilization settings.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Table 4-73. Configurable Policy Rule Elements

Policy Rule Elements	Thresholds, Settings, Definitions
Workload	Enable or disable the demand for memory, CPU, and disk space. Enable or disable the rates for network I/O and datastore I/O, and set the vSphere configuration limit. Configure symptom thresholds for the Workload badge score.
Anomalies	Configure symptom thresholds for the Anomalies badge score.
Faults	Configure symptom thresholds for the Faults badge score.
Capacity Remaining and Time Remaining	Enable or disable the demand and allocation for memory, CPU, and disk space. Enable or disable the rates for network I/O and datastore I/O, and set the vSphere configuration limit. Account for peak times, account for committed projects, which affect the time remaining, and set the provisioning time buffer. Configure thresholds for the Capacity and Time Remaining badge scores.
Stress	Enable or disable the demand for memory and CPU. Enable or disable the rates for network I/O and datastore I/O, and set the vSphere configuration limit. Configure symptom thresholds for the stress badge score.
Reclaimable Capacity	Set the recommended oversize percentage, and the idle and powered off time percentages. Configure symptom thresholds for the Reclaimable Capacity badge score.
Density	Configure symptom thresholds for the Density badge score.
Time	Track the use of objects, and select the maintenance schedule.
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. vRealize Operations Manager treats KPIs differently from other attributes. Threshold violations by a KPI generate different types of alerts from non-KPI attributes. When a KPI violates a threshold, vRealize Operations Manager examines the events that preceded the violation. If it finds enough related information, vRealize Operations Manager captures the set of events that preceded the violation as a fingerprint. If it finds a similar series of events in the future, it can issue a predictive alert warning that the KPI violation is likely to occur.
Alert Definitions	Enable or disable combinations of symptoms and recommendations to identify a condition that classifies as a problem.
Symptom Definitions	Enable or disable test conditions on properties, metrics, or events.

Privileges to Create, Modify, and Prioritize Policies

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

To set the policy priority, on the Active Policies tab, click the policy row and drag it to place it at the desired priority in the list. The priority for the Default Policy is always designated with the letter D.

How Upgrades Affect Your Policies

After you upgrade vRealize Operations Manager from a previous version, you might find newly added or updated default settings of policies such as, new alerts and symptoms. Hence, you must analyze the settings and modify these settings to optimize them for your current environment. If you apply the policies used with a previous version of vRealize Operations Manager, the manually modified policy settings remain unaltered.

Policy Decisions and Objectives

Implementing policy decisions in vRealize Operations Manager is typically the responsibility of the Infrastructure Administrator or the Virtual Infrastructure Administrator, but users who have privileges can also create and modify policies.

You must be aware of the policies established to analyze and monitor the resources in your IT infrastructure.

- As a Virtual Infrastructure Administrator who manages and troubleshoots an IT infrastructure, you must understand how policies associated with objects affect the scores that appear in vRealize Operations Manager, so that you can configure the approved policies based on your company decisions and requirements.
- If you are a Network Operations engineer, you must understand how policies affect the data that vRealize Operations Manager reports on objects, and which policies assigned to objects report alerts and issues.
- If you are the person whose role is to recommend an initial setup for policies, you typically edit and configure the policies in vRealize Operations Manager.
- If your primary role is to assess problems that occur in your environment, but you do not have the responsibility to change the policies, you must still understand how the policies applied to objects affect the data that appears in vRealize Operations Manager. For example, you might need to know which policies apply to objects that are associated with particular alerts.
- If you are a typical application user who receives reports from vRealize Operations Manager, you must have a high-level understanding of the operational policies so that you can understand the reported data values.

Active Policies Tab for Policies

The **Active Policies** tab displays the policies associated with groups of objects. You can manage the active policies for the objects in your environment so that you can have vRealize Operations Manager analyze and display specific data about those objects in dashboards, views, and reports.

How the Active Policies Tab Works

Use the **Active Policies** tab to associate a policy with one or more object groups, and to set the default policy. You can view the locally defined settings for a policy, and the complete list of settings, which includes those that are inherited from the base policies that you select in the Add or Edit Policy workspace. You can assign any policy to be the default policy.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

To display the details for a selected policy, click the split bar to expand the pane. The Details and Related Items tabs and options for the policy appear in the lower pane. On the Related Items tab, you can also apply the selected policy to object groups.

You can use the far right column of the **Active Policies** tab to reorder and therefore reprioritize the policies by dragging them to a new position. However, even though it seems like you can drag a custom policy below the default policy, you cannot. The default policy is always the last policy in the list after the view is refreshed.

How to Prioritize Policies

To set the policy priority, on the Active Policies tab, click the policy row and drag it to place it at the desired priority in the list. The priority for the Default Policy is always designated with the letter D.

Where You Manage the Active Policies

To manage the active policies, in the menu, click **Administration**, and then in the left pane click **Policies**. The **Active Policies** tab appears and lists the policies that are active for the objects in your environment.

Table 4-74. Active Policies Tab Options

Option	Description
Toolbar	Use the toolbar selections to take action on the active policies. ■ Show Association. Opens the **Related Items** tab so that you can associate the policy with groups. ■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to `D`, which gives that policy the highest priority.
Active Policies Tab data grid	vRealize Operations Manager displays the priority and high-level details for the active policies. ■ Priority. Ranking of the priority of the policy. The default policy is marked with a check mark in the Is Default column. ■ Name. Name of the policy as it appears in the Add or Edit Monitoring Policy wizard, and in areas where the policy applies to objects, such as in Custom Groups. ■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects. ■ Groups. Indicates the number of object groups to which the policy is assigned. ■ Affected Objects. Displays the object name, type, and adapter to which the active policy is assigned, and the direct parent group, when applicable. ■ Last Modified. Date and time that the policy was last modified. ■ Modified By. User who last modified the policy settings.

Table 4-74. Active Policies Tab Options (continued)

Option	Description
Active Policies Tab > Details Tab	The Details tab displays the name and description of the policy from which the settings are inherited, the policy priority, who last modified the policy, and the number of object groups associated with the policy. From the Details tab, you can view the settings that are locally defined in your policy, and the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created. ■ Locally Defined Settings. Displays the locally changed policy element settings for each object type in the policy. For example, if you changed the Memory Demand settings in the Cluster Compute Object Stress policy element, you can see the update to your local policy in the list of locally defined settings. ■ Complete Settings Including Inherited. Displays all of the policy element settings for each object type in the policy, including locally changed settings and settings that are inherited. A summary of the enabled and disabled alert definitions, symptom definitions, and attributes appear indicate the number of changes in the policy. The policy element settings include badge score symptom thresholds, and indicate changes made to the Workload, Anomaly, Fault, Capacity and Time Remaining, Stress, Reclaimable Capacity, Density, Usable Capacity, and Time settings. For example, if you changed the Cluster Compute Object Usable Capacity policy element settings, you can see the updates to your local policy in the complete list of settings, and the high availability configuration setting. If you have various adapters installed, such as the vRealize Configuration Manager Adapter, you will also see specific policy elements for the adapter. For example, for vRealize Configuration Manager you will see the Compliance policy element setting and badge score symptom threshold.

Active Policies Tab > Related Objects Tab	Summarizes the related groups and objects, and details about the selected object group and objects.
	■ Groups. Displays the groups of objects associated with the selected active policy, and provides options to add and release an association.
	  ■ Add Association. Opens the Apply the policy to groups dialog box where you select object groups to associate with the selected policy.
	  ■ Release Association. Opens a confirmation dialog box to confirm the release of the object group that is associated with the selected policy.
	  ■ Data grid. Displays the groups assigned to this policy, the object types associated with the group, and the number of objects in the group.
	  ■ Details for the selected object group. Displays the object group name, type, and number of members associated with the selected policy, and the type of association with the policy. An object group can have a direct association with a policy, and inherited policy associations based on the base policies that you selected when you created a local policy. For example, if the Base Settings policy appears in the list, with an inherited association, the Base Settings policy was included in the base policies selected when this policy was created.
	■ Affected Objects. Displays the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it appears in this data grid.

Policy Library Tab for Policies

The **Policy Library** tab displays the base settings, default policy, and other best practice policies that vRealize Operations Manager includes. You can use the library policies to create your own

policies. The policy library includes all the configurable settings for the policy elements, such as workload, anomaly, faults, capacity and time remaining, stress, reclaimable capacity, density, usable capacity, and time.

How the Policy Library Works

Use the options on the **Policy Library** tab to create your own policy from an existing policy, or to override the settings from an existing policy so that you can apply the new settings to groups of objects. You can also import and export a policy.

To display the details for a selected policy, click the split bar to expand the pane. The Details and Related Items tabs and options for the policy appear in the lower pane. On the Related Items tab, you can also apply the selected policy to object groups.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can also apply the policy to object groups. To update the policy association with an object group, the role assigned to your user account must have the Manage Association permission enabled for policy management.

Where You Manage the Policy Library

To manage the policy library, in the menu, click **Administration**, and then in the left pane click **Policies**. The **Policy Library** tab appears and lists the policies available to use for your environment.

Table 4-75. Policy Library Tab Options

Option	Description
Toolbar	Use the toolbar selections to take action in the policy library. ■ Add New Policy. Create a policy from an existing policy. ■ Edit Selected Policy. Customize the policy so that you can override settings for vRealize Operations Manager to analyze and report data about the associated objects. ■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to `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. ■ Delete Selected Policy. Remove a policy from the list.
Policy Library Tab 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 Monitoring Policy wizard, and in areas where the policy applies to objects, such as in Custom Groups. ■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects. ■ Last Modified. Date and time that the policy was last modified. ■ Modified By. User who last modified the policy settings.

Table 4-75. Policy Library Tab Options (continued)

Option	Description
Policy Library Tab > Details Tab	The Details tab displays the name and description of the policy from which the settings are inherited, the policy priority, who last modified the policy, and the number of object groups associated with the policy. From the Details tab, you can view the settings that are locally defined in your policy, and the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created. ■ Locally Defined Settings. Displays the locally changed policy element settings for each object type in the policy. For example, if you changed the Memory Demand settings in the Cluster Compute Object Stress policy element, you can see the update to your local policy in the list of locally defined settings. ■ Complete Settings Including Inherited. Displays all of the policy element settings for each object type in the policy, including locally changed settings and settings that are inherited. A summary of the enabled and disabled alert definitions, symptom definitions, and attributes appear indicate the number of changes in the policy. The policy element settings include badge score symptom thresholds, and indicate changes made to the Workload, Anomaly, Fault, Capacity and Time Remaining, Stress, Reclaimable Capacity, Density, Usable Capacity, and Time settings. For example, if you changed the Cluster Compute Object Usable Capacity policy element settings, you can see the updates to your local policy in the complete list of settings, and the high availability configuration setting. If you have various adapters installed, such as the vRealize Configuration Manager Adapter, you will also see specific policy elements for the adapter. For example, for vRealize Configuration Manager you will see the Compliance policy element setting and badge score symptom threshold.

Related Objects Tab	Summarizes the related groups and objects, and details about the selected object group and objects.
	■ Groups. Displays the groups of objects associated with the selected active policy, and provides options to add and release an association.
	  ■ Add Association. Opens the Apply the policy to groups dialog box where you select object groups to associate with the selected policy.
	  ■ Release Association. Opens a confirmation dialog box to confirm the release of the object group that is associated with the selected policy.
	  ■ Data grid. Displays the groups assigned to this policy, the object types associated with the group, and the number of objects in the group.
	  ■ Details for the selected object group. Displays the object group name, type, and number of members associated with the selected policy, and the type of association with the policy. An object group can have a direct association with a policy, and inherited policy associations based on the base policies that you selected when you created a local policy. For example, if the Base Settings policy appears in the list, with an inherited association, the Base Settings policy was included in the base policies selected when this policy was created.
	■ Affected Objects. Displays the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it appears in this data grid.

Operational Policies

Determine how to have vRealize Operations Manager monitor your objects, and how to notify you about problems that occur with those objects.

vRealize Operations Manager Administrators assign policies to object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Enable and disable alerts.
- Control data collections by persisting or not persisting metrics on the objects in your environment.
- Configure the product analytics and thresholds.
- Monitor objects and applications at different service levels.
- Prioritize policies so that the most important rules override the defaults.
- Understand the rules that affect the analytics.
- Understand which policies apply to object groups.

vRealize Operations Manager includes a library of built-in active policies that are already defined for your use. vRealize Operations Manager applies these policies in priority order.

When you apply a policy to an object group, vRealize Operations Manager collects data from the objects in the object group based on the thresholds, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that are enabled in the policy.

The following examples of policies might exist for a typical IT environment.

- Maintenance: Optimized for ongoing monitoring, with no thresholds or alerts.
- Critical Production: Production environment ready, optimized for performance with sensitive alerting.
- Important Production: Production environment ready, optimized for performance with medium alerting.
- Batch Workloads: Optimized to process jobs.
- Test, Staging, and QA: Less critical settings, fewer alerts.
- Development: Less critical settings, no alerts.
- Low Priority: Ensures efficient use of resources.
- Default Policy: Default system settings.

User Scenario: Create an Operational Policy for Production vCenter Server Datastore Objects

As a Virtual Infrastructure Administrator, you manage the policies used for vRealize Operations Manager to analyze objects in your environment, collect data from those objects, and display that data in dashboards, views, and reports. Your IT staff added new datastore objects to your

environment, and your responsibility is to ensure that the new datastore objects adhere to the policy requirements from the VP of Infrastructure for your test and production environments.

In this scenario, you create a policy to have vRealize Operations Manager monitor the disk space use of your production datastore objects. You create a group type and custom object group for the datastore objects, and apply your policy to your object group. After vRealize Operations Manager collects data from the datastore objects in your environment according to the settings in your policy, you view the collected data and any potential alerts in the dashboards to confirm whether the disk space use is in compliance for your datastore objects.

Prerequisites

- Understand the purpose of using a policy. See Policies .
- Verify that your vRealize Operations Manager instance is working properly.
- Verify that one or more custom object groups and group types exist in your vRealize Operations Manager instance. See Managing Custom Object Groups in VMware vRealize Operations Manager.
- Verify that your vRealize Operations Manager instance includes the default policy and one or more other policies. See Default Policy in vRealize Operations Manager.
- Understand the sections and elements in the default policy, such as the attributes, alert and symptom definitions, and how the policy inherits settings from the base policies that you select. See Policy Workspace in vRealize Operations Manager.
- Understand the analysis settings in the default policy, such as capacity remaining and stress on hosts and virtual machines, and the actions used to override the settings inherited from the base policies. See Analysis Settings Details.

Procedure

1 Create a Group Type for Your Datastore Objects

   Create a group type so that you can categorize your Datastore objects.

2 Create an Object Group for Your Datastore Objects

   Create an object group to organize the Datastore objects in your environment as a single object group.

3 Create Your Policy and Select a Base Policy

   Create your policy, and select the base policies to use to override the settings for your new policy.

4 Override the Analysis Settings for the Datastore Objects

   Display and override the analysis settings for the Datastore objects that your new policy will monitor.

5 Enable Disk Space Attributes for Datastore Objects

   Enable the attributes for vRealize Operations Manager to monitor the disk space of your production datastore objects.

6 Override Alert and Symptom Definitions for Datastore Objects

Override the alert and symptom definitions for Datastore objects.

7 Apply Your Datastore Policy to Your Datastore Objects Group

Apply the policy to your new group of Datastore objects to have vRealize Operations Manager monitor them to ensure that the disk space levels of these objects adhere to the settings in your policies to support the service level agreements and business priorities that are established for your environment.

8 Create a Dashboard for Disk Use of Your Datastore Objects

Create a dashboard so that you can monitor the disk use of your Datastore objects, and be alerted to any potential problems.

Results

You created a policy to apply to your new production Datastore objects so that you can have vRealize Operations Manager monitor them to ensure that the disk space levels of these objects adhere to the settings in your policies to support the service level agreements and business priorities that are established for your environment. vRealize Operations Manager uses the settings in your new policy to display the disk use for your Datastore objects in dashboards, views, and reports, and to enforce the service levels during data collections.

What to do next

After you finish this scenario, you must wait for vRealize Operations Manager to collect data from the objects in your environment. Then view the disk use of your Datastore objects.

Create a Group Type for Your Datastore Objects

Create a group type so that you can categorize your Datastore objects.

In this step, you create a group type so that you can apply it to the new custom object group that you create to organize your vCenter Server Datastore objects.

Prerequisites

Verify that you understand the context of this scenario. See User Scenario: Create an Operational Policy for Production vCenter Server Datastore Objects.

Procedure

1 In the menu, click **Administration**, and then in the left pane click **Configuration > Group Types**.

2 Click the **Add Group Type** icon to add a new group type, type **`Production_Datastores`**, and click **OK**.

The new group type appears in the list of group types.

What to do next

Create an object group so that you can organize the Datastore objects in your environment as a single object group.

Create an Object Group for Your Datastore Objects

Create an object group to organize the Datastore objects in your environment as a single object group.

In this step, you create a new object group to organize your Datastore objects so that you can apply the policy that you create to the object group.

Prerequisites

Create an object type. See Create a Group Type for Your Datastore Objects.

Procedure

1. In the menu, click **Environment**, and click **Custom Groups**.
2. On the **Groups** tab, click the plus sign to add a new group, and enter a name for the object group.
3. From the **Group Type** drop-down menu, select your new group type.
4. From the **Policy** drop-down menu, select the Default Policy for now.

   To have vRealize Operations Manager identify new Datastore objects that are added to your environment, you select the **Keep group membership up to date** check box to make this group dynamic and keep it updated.
5. In the Define membership criteria pane, select the **vCenter Adapter > Datastore** object type from the drop-down menu.
6. Click in the **Pick a property** text box, and select **Disk Space > Template > Virtual Machine used (GB)**.
7. In the adjacent text box, click the drop-down arrow and select **is less than**.
8. In the **Property value** text box, type **10**.

   vRealize Operations Manager uses this criteria to monitor Datastore objects in this group, and to report when the Datastore objects have less than 10 GB of space remaining.
9. In the Objects to always include pane, select the object group that you created for your Datastore objects, click **Add** to move the group to the selected pane, and select the object group check box.

   In the Objects to always exclude pane, do not select objects to exclude.
10. Click **OK** to save your new group.

What to do next

Create your policy, and select the base policies to use to override the settings for your new policy.

Create Your Policy and Select a Base Policy

Create your policy, and select the base policies to use to override the settings for your new policy.

In this step, you create a policy for vRealize Operations Manager to analyze and monitor your Datastore objects, and select the policies from which to inherit and override the settings for your new policy.

Prerequisites

Create a custom object group for your Datastore objects. See Create an Object Group for Your Datastore Objects.

Procedure

1. Access the Policies area to create your policy.
   a. In the menu, click **Administration**, and then in the left pane click **Policies**.

      The **Active Policies** and **Policy Library** tabs appear.

   b. Click the **Policy Library** tab, and click the plus sign to add a policy.

   c. In the Getting Started policy workspace, enter a name and description for the policy.

   d. In the Start with area, select **Default Policy** to inherit settings from a base policy.
2. Select the base policies, object, and policy to use to override the settings for your new policy.
   a. In the policy workspace, click **Select Base Policies**.

   b. To view the current policy configuration for your Datastore objects, click the **Show changes for** drop-down menu, click **vCenter Adapter - Datastore**, and click the **Show object type** filter.

      The Datastore policy configuration appears in the right pane.

What to do next

Display and override the analysis settings for the Datastore objects that your new policy will monitor.

Override the Analysis Settings for the Datastore Objects

Display and override the analysis settings for the Datastore objects that your new policy will monitor.

In this step, you override the capacity remaining and time remaining settings for your new policy, and override the capacity score symptom thresholds so that vRealize Operations Manager triggers an alert and notifies you of potential problems with the capacity of your Datastore objects.

Prerequisites

Create your policy and select the base policies to inherit and override the settings for your new policy. See Create Your Policy and Select a Base Policy.

Procedure

1. In the policy workspace, click **Analysis Settings**.
2. Click the **Show changes for** drop-down menu, click **vCenter Adapter - Datastore**, and click the **Show object type** filter.

   The vCenter Adapter - Datastore object type appears in the Object types list, and the analysis settings for Datastore objects appear in the right pane. The policy elements include thresholds and settings for all of the analysis capabilities, such as Workload, Stress, Usable Capacity, and so on.
3. Click the policy element override button for the Capacity Remaining and Time Remaining element to turn on this policy element.

   The button changes to a check mark, and the policy element becomes active so that you can override the settings.
4. Click and drag the settings on the Capacity Score Symptom Threshold slider to 10% for warning (red), 15% for caution (orange), and 20% for normal (green).

   When these thresholds are violated for the Datastore objects in your environment, vRealize Operations Manager triggers an alert and notifies you of a potential problem with the capacity of your Datastore objects.
5. Click the policy element override button for the Usable Capacity element to turn on this policy element, click the arrow to expand the policy element view, and select the **Use High Availability (HA) Configuration** check box.

   When you use High Availability, you ensure that vRealize Operations Manager provides enough resources for your Datastore objects to handle throughput and potential loss of data.

What to do next

Enable the disk space attributes for datastore objects.

Enable Disk Space Attributes for Datastore Objects

Enable the attributes for vRealize Operations Manager to monitor the disk space of your production datastore objects.

In this step, you enable vRealize Operations Manager to monitor and collect the disk space properties attribute from the Datastore objects in your environment.

Prerequisites

Override the analysis settings for your Datastore objects. See Override the Analysis Settings for the Datastore Objects.

Procedure

1. In the policy workspace, click **Override Attributes**.
2. From the Object Type drop-down menu, select **vCenter Adapter > Datastore**.

   vRealize Operations Manager filters the list and displays only the attributes that apply to Datastore objects.
3. Click the **Attribute Type** drop-down menu, select **Property**, and deselect the other attributes.
4. Enter **space** in the **Search** text box, and click the search button.

   vRealize Operations Manager filters the list and displays only the disk space properties associated with Datastore objects.
5. For the **Disk Space|Template|Virtual Machine used (GB)** property attribute, click the **State** drop-down menu, and click **Local**.

   When this attribute is enabled in your local policy, vRealize Operations Manager collects this disk space properties attribute from Datastore objects in your environment.

What to do next

Override the alert symptom definitions for Datastore objects.

Override Alert and Symptom Definitions for Datastore Objects

Override the alert and symptom definitions for Datastore objects.

In this step, you override the alert and symptom definitions so that vRealize Operations Manager uses trigger an alert notification during data collections when the disk space for your Datastore objects begins to run out.

Prerequisites

Enable vRealize Operations Manager to monitor and collect the disk space properties attribute from the Datastore objects in your environment. See Enable Disk Space Attributes for Datastore Objects.

Procedure

1. In the policy workspace, click **Alert / Symptom Definitions**.
2. In the Alert Definitions pane, from the Object Type drop-down menu, select **vCenter Adapter > Datastore**.
3. Enter **space** in the **Search** text box, and click the search button.

4 For the alert definition named `Datastore is running out of disk space`, click the **State** drop-down menu and click **Local**.

   When this alert definition is enabled in your local policy, vRealize Operations Manager uses it to trigger an alert notification during data collections when the disk space for your Datastore objects begins to run out.

5 In the Symptom Definitions pane, from the Object Type drop-down menu, select **vCenter Adapter > Datastore**.

6 Enter ***space*** in the **Search** text box, and click the search button.

7 To enable the critical, immediate, and warning symptom definitions for the space use on datastore objects, click **Actions**, and click **Select All**, then set the thresholds.

Table 4-76. Symptom Definitions Threshold Settings

Selection	Setting
Datastore space use reaching critical limit	>90
Datastore space use reaching immediate limit	>85
Datastore space use reaching warning limit	>80

What to do next

Apply your policy to your Datastore objects.

Apply Your Datastore Policy to Your Datastore Objects Group

Apply the policy to your new group of Datastore objects to have vRealize Operations Manager monitor them to ensure that the disk space levels of these objects adhere to the settings in your policies to support the service level agreements and business priorities that are established for your environment.

In this step, you apply your new policy to production Datastore objects so that vRealize Operations Manager monitors them to ensure adequate disk space levels of these objects.

Prerequisites

Override the alert and symptom definitions for Datastore objects. See Override Alert and Symptom Definitions for Datastore Objects.

Procedure

1 In the policy workspace, click **Apply Policy to Groups**, and select the new object group that you created for your Datastore objects.

2 Click **Save** to save your new policy settings.

Results

vRealize Operations Manager uses the settings in your new policy to display the disk use for your Datastore objects in dashboards, views, and reports, and to enforce the service levels during data collections

What to do next

Create a new dashboard to view the disk use of your Datastore objects.

Create a Dashboard for Disk Use of Your Datastore Objects

Create a dashboard so that you can monitor the disk use of your Datastore objects, and be alerted to any potential problems.

In this step, you create a new dashboard, add widgets to your new dashboard, and configure the widgets so that you can monitor your production datastore objects.

Prerequisites

Apply the policy to your new group of Datastore objects. See Apply Your Datastore Policy to Your Datastore Objects Group.

Procedure

1. In the menu, click **Dashboards**, and then click **Actions > Create Dashboard**.
2. Configure your new dashboard.
   a. In the Dashboard Configuration pane of the New Dashboard workspace, enter the name **`Production Datastores`** for the new dashboard.
   b. For Is default, select **Yes**.
3. Add widgets to your new dashboard.
   a. In the workspace, click **Widget List**.
   b. From the list of widgets, click the **Object List** widget, and drag it to the right pane.
   c. Click the **Capacity** widget, and drag it to the right pane.
   d. Click the **Time Remaining** widget, and drag it to the right pane.
   e. Click the **Alert List** widget, and drag it to the right pane.
4. Configure the widget interactions.
   a. In the workspace, click **Widget Interactions**.
   b. For the Object List widget interactions, click the drop-down menu for the Selected Objects and Selected Alerts, and clear the selections.
   c. For the Alert List widget interaction, click the drop-down and select **Object List**.
   d. For the Capacity widget interaction, click the drop-down and select **Object List**.

e For the Time Remaining widget interaction, click the drop-down and select **Object List**.

f Click **Apply Interactions**.

5 Configure the Object List widget.

a On the Object List widget, click the pencil.

b For Refresh Content, select **On**.

c For Refresh Interval, click the arrows and select **30** seconds.

d For Mode, select **Parent**.

e For Auto Select First Row, select **Off**.

f In the lower pane, click the plus sign to expand the list of tags, expand **Production Datastores**, select **Production Datastores (n)**, and click **OK**.

The objects in your Production Datastores object group appears in the Object List widget.

6 Configure the Capacity widget.

a On the Capacity widget, click the pencil.

b For Refresh Content, select **On**.

c For Refresh Interval, click the arrows and select **30** seconds.

d For Self Provider, select **On**.

e For Selected Object, in the **Search** text box, enter **group**, and select the **Production Datastores** group from the list.

The Production Datastores group appears in the **Selected Object** text box.

f Click **OK**.

The Capacity widget displays a score and a graph to indicate the remaining compute objects as a percentage of the total consumer capacity.

7 Configure the Time Remaining widget.

a On the Time Remaining widget, click the pencil.

The Time Remaining widget displays the amount of time that remains until the object resources are consumed.

b For Refresh Content, select **On**.

The Time Remaining widget displays the amount of time that remains until the object resources are consumed.

c For Refresh Interval, click the arrows and select **30** seconds.

d For Self Provider, select **On**.

e For Selected Object, in the **Search** text box, enter **group**, and select the **Production Datastores** group from the list.

The Production Datastores group appears in the **Selected Object** text box.

f Click **OK**.

The Time Remaining widget displays a score and a graph to indicate the amount of time that remains until the object resources are consumed.

8 Configure the Alert List widget.

a On the Alert List widget, click the pencil.

b For Refresh Content, select **On**.

c For Refresh Interval, click the arrows and select **30** seconds.

d For Selected Object, in the **Search** text box, enter **group**, and select the **Production Datastores** group from the list.

The Production Datastores group appears in the **Selected Object** text box.

e In the lower pane, click the plus sign to expand the list of tags, expand **Production Datastores**, select **Production Datastores (n)**, and click **OK**.

The alert list widget displays the alerts that are configured for your objects. You created a dashboard to monitor disk space of your production datastore objects. After vRealize Operations Manager analyzes and collects data from the objects in your Production Datastores object group, you can view the results in your new dashboard.

Results

You created and applied a policy to your production datastore objects to have vRealize Operations Manager monitor those objects during data collections so that you can monitor and enforce the service levels for your environment. vRealize Operations Manager uses the settings in your new policy to display information about the capacity, time remaining, and potential alerts for your Datastore objects. With your new policy in place, you can ensure that the disk space levels for your production datastore objects adhere to the policies established for your production environment.

Types of Policies

There are three types of policies such as default policies, custom policies, and policies that are offered with vRealize Operations Manager.

Custom Policies

You can customize the default policy and base policies included with vRealize Operations Manager for your own environment. You can then apply your custom policy to groups of objects, such as the objects in a cluster, or virtual machines and hosts, or to a group that you create to include unique objects and specific criteria.

You must be familiar with the policies so that you can understand the data that appears in the user interface, because policies drive the results that appear in the vRealize Operations Manager dashboards, views, and reports.

To determine how to customize operational policies and apply them to your environment, you must plan ahead. For example:

- Must you track CPU allocation? If you overallocate CPU, what percentage must you apply to your production and test objects?
- Will you overallocate memory or storage? If you use High Availability, what buffers must you use?
- How do you classify your logically defined workloads, such as production clusters, test or development clusters, and clusters used for batch workloads? Or, do you include all clusters in a single workload?
- How do you capture peak use times or spikes in system activity? In some cases, you might need to reduce alerts so that they are meaningful when you apply policies.

When you have privileges applied to your user account through the roles assigned, you can create and modify policies, and apply them to objects. For example:

- Create a policy from an existing base policy, inherit the base policy settings, then override specific settings to analyze and monitor your objects.
- Use policies to analyze and monitor vCenter Server objects and non-vCenter Server objects.
- Set custom thresholds for analysis settings on all object types to have vRealize Operations Manager report on workload, anomalies, faults, capacity, stress, and so on.
- Enable specific attributes for collection, including metrics, properties, and super metrics.
- Enable or disable alert definitions and symptom definitions in your custom policy settings.
- Apply the custom policy to object groups.

When you use an existing policy to create a custom policy, you override the policy settings to meet your own needs. You set the allocation and demand, the overcommit ratios for CPU and memory, and the thresholds for capacity risk and buffers. To allocate and configure what your environment is actually using, you use the allocation model and the demand model together. Depending on the type of environment you monitor, such as a production environment versus a test or development environment, whether you over allocate at all and by how much depends on the workloads and environment to which the policy applies. You might be more conservative with the level of allocation in your test environment and less conservative in your production environment.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The

default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Your policies are unique to your environment. Because policies direct vRealize Operations Manager to monitor the objects in your environment, they are read-only and do not alter the state of your objects. For this reason, you can override the policy settings to fine-tune them until vRealize Operations Manager displays the results that are meaningful and that affect for your environment. For example, you can adjust the capacity buffer settings in your policy, and then view the data that appears in the dashboards to see the effect of the policy settings.

User Scenario: Create a Custom Operational Policy for a vSphere Production Environment

As a system administrator of vRealize Operations Manager, you are responsible for ensuring that the objects in your vSphere environment conform to specific policies. You must ensure that your objects have enough memory and CPU to support your Test, Development, and Production environments.

Large IT environments might include four to six production environments that are organized according to object types, with a minor policy applied to each area. These large environments typically include a default policy, a single production policy that applies to the entire environment, and individual policies for dedicated areas.

You typically apply a default policy to most of the objects in your environment. To have vRealize Operations Manager monitor and analyze dedicated groups of objects, you create a separate policy for each object group, and make only minor changes in the settings for that policy. For example, you might apply a default operational policy for all of the objects in your vSphere production environment, but you also need to closely track the health and risk of virtual SQL Server instances, including their capacity levels. To have vRealize Operations Manager analyze only the virtual SQL Server instances, and to monitor them, you create a separate, dedicated policy and apply that policy to that group of objects. The settings in the policy that you create to monitor the virtual SQL Server instances differs only slightly from the main production policy.

This scenario shows you how to use multiple policies to analyze and monitor specific objects, so that you can manage them to ensure continuous operation. In this scenario, your vSphere production environment is one part of your overall production environment. You must create a custom operational policy to monitor the virtual SQL Server objects in your vSphere production environment.

Prerequisites

- Understand the purpose of using a policy. See Policies .
- Verify that your vRealize Operations Manager instance is working properly.
- Verify that your vRealize Operations Manager instance includes the Default Policy and one or more other policies. See Default Policy in vRealize Operations Manager.

- Understand the sections and elements in the policy, such as the attributes, alert and symptom definitions, and how the policy inherits settings from the base policies that you select. See Policy Workspace in vRealize Operations Manager.
- Understand the analysis settings in the policy, such as capacity remaining and stress on hosts and virtual machines, and the actions used to override the settings inherited from the base policies. See Analysis Settings Details.

Procedure

1. Determine the vSphere Operational Requirements

   You must continuously monitor the capacity levels of your virtual SQL Server machines, and have vRealize Operations Manager notify you about any degradation in the performance of these objects. You want vRealize Operations Manager to notify you 60 days before these objects begin to experience problems with their capacity levels.

2. Create a Policy to Meet vSphere Operational Needs

   You create an operational policy for your virtual SQL Server instances, where only these settings differ from the main production policy. In this policy, you change the memory and CPU settings for specific objects. You then configure vRealize Operations Manager to send alerts to you when the performance degrades on your virtual SQL Servers.

3. Configure the Custom Policy Settings to Analyze and Report on vSphere Objects

   You use different policy requirements for your Development, Test, and Production environments so that you can configure the specific policy settings for vRealize Operations Manager to analyze and report on your objects, including your virtual SQL Servers.

4. Apply the Custom Policy to vSphere Object Groups

   You create an object group type to categorize your virtual SQL Server machines. Then you create an object group that contains your virtual SQL Server machines, and apply your custom policy to this group of SQL Server virtual machine objects.

What to do next

After you finish this scenario, you must wait for vRealize Operations Manager to collect data from the objects in your environment. When a violation of the policy thresholds occur, vRealize Operations Manager sends an alert to notify you of the problem. If you continuously monitor the state of your objects, you are always aware of the state of the objects in your environment, and do not need to wait for vRealize Operations Manager to send alerts.

Create a custom dashboard so that you can monitor the virtual SQL Server objects and address problems that occur. See Dashboards.

Determine the vSphere Operational Requirements

You must continuously monitor the capacity levels of your virtual SQL Server machines, and have vRealize Operations Manager notify you about any degradation in the performance of these objects. You want vRealize Operations Manager to notify you 60 days before these objects begin to experience problems with their capacity levels.

Your VP of Infrastructure has defined a default operational policy and a main production policy for all of the objects in your production environment, and your IT Director has applied these policies to your production environments. Although the main production policy handles the operational monitoring needs for most of your objects, your manager requires that you be notified about any degradation in the performance of your production virtual SQL Server machines. You have vRealize Operations Manager continuously monitor the capacity levels of your virtual SQL Servers so that you can address problems that occur. You have vRealize Operations Manager notify you 60 days before your virtual SQL Servers begin to experience problems with their capacity levels.

Your IT department divided objects into dedicated groups that support the Development, Test, and Production areas. You must use vRealize Operations Manager to continually track and assess the health and risk of the objects in each of these areas.

In this scenario, you create an operational management policy to analyze, monitor, and troubleshoot your objects. You then monitor the results in custom dashboards.

You must first determine the vSphere operational requirements so that you can understand the analysis settings required for your policy. You can then create a policy to monitor your virtual SQL Server objects, and configure the custom policy to include minor differences in the settings for the main production policy.

When you create the custom policy to analyze and monitor your virtual SQL Servers, you configure the analysis settings so that vRealize Operations Manager analyzes specific objects and report the results in the dashboards. You then apply the policy to groups of virtual SQL Server objects.

Prerequisites

Verify that the following conditions are met:

- You understand the context of this scenario. See User Scenario: Create a Custom Operational Policy for a vSphere Production Environment.
- A default policy and a main production policy are in effect for all of the objects in your vSphere production environment.

Procedure

1. Determine the operational requirements for your vSphere production environment.

   In this scenario, the following requirements will be applied to the environment.

2. Develop a plan to create a custom operational policy that meets the requirements to analyze and monitor the objects in your environment.

   a. Ensure that virtual SQL Servers continuously have adequate memory and CPU capacity.

   b. Ensure that you do not overcommit memory on your production virtual SQL Servers.

c Overcommit only a small percentage of the CPUs on your SQL Servers.

In this scenario, you set the value to 2. In some production environments, a typical value might be 4.

d Ensure that vRealize Operations Manager alerts you if the capacity of your virtual SQL Servers drops below the defined thresholds.

e Set the Co-Stop value on your production virtual SQL Servers to an acceptable level so that the SQL Servers do not experience delay because of CPU scheduling contention.

f Determine whether to overcommit compute resources for certain ratios.

Results

After you plan the custom policy requirements, you can implement the policy.

What to do next

Create an operational policy for your virtual SQL Server instances.

Create a Policy to Meet vSphere Operational Needs

You create an operational policy for your virtual SQL Server instances, where only these settings differ from the main production policy. In this policy, you change the memory and CPU settings for specific objects. You then configure vRealize Operations Manager to send alerts to you when the performance degrades on your virtual SQL Servers.

In this procedure, you create a dedicated policy for a subset of virtual SQL Server objects, and change settings for the memory and CPU capacity for your virtual SQL Server instances. At this point in the scenario, your custom policy has only minor differences from the production policy.

The difference between the main production policy and your virtual SQL Server policy is in the overcommitment of compute resources. For the SQL Server policy, you do not overcommit compute resources. You have the SQL server policy inherit most of the settings from your overall production policy, except that you change the capacity settings that apply directly to the virtual SQL servers.

After you apply the main production policy to your entire production environment, you create the dedicated policy, have it inherit settings from the main policy, and make minor changes to settings in the dedicated policy to adjust the capacity levels for your virtual SQL Servers.

To create this policy, you choose a cluster that contains the data center and the vCenter Server that uses this policy. You make minor changes for all the objects, including the cluster, data center, host system, resource pools, and the virtual machine resource containers.

Prerequisites

Verify that the following conditions are met:

- You know the vSphere operational requirements. See Determine the vSphere Operational Requirements.
- A default policy is in effect for your entire production environment of vSphere objects.

Procedure

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

The **Active Policies** tab displays the current policies in effect.

2 Click the **Policy Library** tab, and click the **Add New Policy** icon to add a custom policy.

3 In the workspace navigation pane, click **Getting Started** and define the basic information for the policy.

a In the **Name** text box, enter **vSphere Production Virtual SQL Servers**.

b In the **Description** text box, enter **Analyze capacity of virtual SQL Servers**.

c To start with a base policy, select **Default Policy** from the **Start with** drop-down menu.

4 View the policy configuration settings.

a In the policy workspace, click **Select Base Policies**.

b To view the policy configuration for virtual machine objects, click the **Show changes for** drop-down menu, click **vCenter Adapter - Virtual Machine**, and click the **Show object type** filter.

The Virtual Machine policy configuration appears in the right pane.

c To view the inherited settings, in the Policy Preview pane, click **Configuration inherited from base policy**.

5 In the workspace navigation, click **Analysis Settings**.

6 In the workspace navigation, add the following object types to the list so that you can change their settings.

a Click the drop-down arrow, click **vCenter Adapter - Cluster Compute Resource**, and click the filter.

b Click the drop-down arrow, click **vCenter Adapter - Data Center**, and click the filter.

c Click the drop-down arrow, click **vCenter Adapter - Host System**, and click the filter.

d Click the drop-down arrow, click **vCenter Adapter - Resource Pool**, and click the filter.

e Click the drop-down arrow, click **vCenter Adapter - Virtual Machine**, and click the filter.

The analysis settings for these object types appear in the right pane.

7 On the Cluster Compute Resource bar, click the double arrows to expand the list of analysis settings.

8 Locate **Capacity Remaining Time Remaining** and click the lock button to enable changes.

9 In the resource table, set the overcommit for Memory Allocation value to **0** so that vRealize Operations Manager does not overcommit these objects for your SQL Server policy.

10 In the resource table, set the overcommit ratio for CPU Allocation to **2** so that vRealize Operations Manager overcommits a 2:1 ratio for CPU allocation on each SQL Server.

11 Repeat Step 7 through Step 10 for each object type that you added to the right pane.

12 Click **Save**.

Results

You created a policy and made minor changes to settings so that vRealize Operations Manager can analyze and report on your SQL Server objects.

What to do next

Configure the alert definitions and symptom definitions for your SQL Server policy. You apply the policy to your SQL Server object groups.

Configure the Custom Policy Settings to Analyze and Report on vSphere Objects

You use different policy requirements for your Development, Test, and Production environments so that you can configure the specific policy settings for vRealize Operations Manager to analyze and report on your objects, including your virtual SQL Servers.

This scenario presents several typical cases where you might be required to differentiate between the policy requirements for Development, Test, and Production environments.

- For your Development and Test environments, you might not be concerned if the objects in these environments experience network redundancy loss, but you do care when the objects fail. In this case, you locate the Physical NIC link state alert definition, double-click the state, and set it to Disabled.
- For a Test environment, you might not be concerned if your virtual machines demand more memory and CPU capacity than what is actually configured, because workloads can vary in test environments.
- For a Production environment, your virtual machines might require more memory than you have configured, which might cause a problem with the performance and reliability of your production environment.

In this procedure, you override the symptom definition threshold value for the Co-Stop performance of your virtual machines.

Prerequisites

Verify that the following conditions are met:

- You created a custom policy for your virtual SQL Servers. See Create a Policy to Meet vSphere Operational Needs.
- You understand the Co-Stop CPU performance metric for virtual machines. This metric represents the percentage of time that a virtual machine is ready to run, but experiences delay because of co-virtual CPU scheduling contention. Co-Stop is one of several performance metrics for virtual machines that also include Run, Wait, and Ready.
- The alert definition named Virtual machine has high CPU contention caused by Co-Stop, exists.

- Symptom definitions exist to track the critical, immediate, and warning levels of CPU Co-Stop on the virtual machines. For example, the critical level for virtual machine CPUs that experience contention more than 15% of the time is set to 15% by default, as measured by the Co-Stop metric. The default threshold level for Immediate is 10%, and for warning is 5%. However, in your production policy for your production virtual machines, you manage the critical level at 3%.

Procedure

1 On the **Policy Library** tab, locate your vSphere Production Virtual SQL Servers policy, and click the pencil to edit the policy.

  The Edit Monitoring Policy workspace appears.

2 In the workspace, click **Override Alert / Symptom Definitions**.

3 On the Alert Definitions pane, enable the Co-Stop alert definition to notify you about high CPU contention on your virtual machines.

  a In the Object Type drop-down menu, select **vCenter Adapter** and **Virtual Machine**.

  b In the **Search** text box, enter **`stop`** to display only the alert definitions that relate to the Co-Stop performance metric for virtual machines.

  c For the Alert definition named `Virtual machine has high CPU contention caused by Co-Stop`, click the **State** drop-down menu and click **Enabled**.

4 In the Symptom Definitions pane, modify the critical Co-Stop level for virtual machines so that vRealize Operations Manager triggers an alert based on the threshold level defined for this symptom.

  a In the Object Type drop-down menu, click **vCenter Adapter** and **Virtual Machine**.

  b In the **Search** text box, enter **`stop`** to display the symptom definitions that apply to the Co-Stop performance metric for virtual machines.

  c For the symptom definition named `Virtual Machine CPU Co-stop is at Critical level`, click the **State** drop-down menu and click **Enabled**.

  d Click the **Condition** drop-down menu, and click **Override**.

    For a production policy, a typical critical threshold value is **`>3`**. For a development or test environment policy, a typical critical threshold value is **`>10`**.

  e In the Override Symptom Definition Threshold dialog box, enter **`>3`** to change the threshold value, and click **Apply.**

5 Modify the immediate Co-Stop level for virtual machines.

   a For the symptom definition named `Virtual Machine CPU Co-stop is at Immediate level`, click the **State** drop-down menu and click **Enabled**.

   b Click the **Condition** drop-down menu, and click **Override**.

   c In the Override Symptom Definition Threshold dialog box, enter **>2** to change the threshold value, and click **Apply.**

6 Modify the warning Co-Stop level for virtual machines.

   a For the symptom definition named `Virtual Machine CPU Co-stop is at Warning level`, click the **State** drop-down menu and click **Enabled**.

   b Click the **Condition** drop-down menu, and click **Override**.

   c In the Override Symptom Definition Threshold dialog box, enter **>1** to change the threshold value, and click **Apply**.

7 Click **Save** to save your policy.

Results

You changed the Co-Stop CPU performance metric for virtual machines to minimize the delay on your SQL Server virtual machines because of CPU scheduling contention.

What to do next

Create a group type to use to categorize your group of virtual SQL Servers, create an object group that contains your virtual SQL Servers, and apply the policy to your object group.

Apply the Custom Policy to vSphere Object Groups

You create an object group type to categorize your virtual SQL Server machines. Then you create an object group that contains your virtual SQL Server machines, and apply your custom policy to this group of SQL Server virtual machine objects.

To have vRealize Operations Manager analyze your SQL Server machines according to the performance criteria in your custom policy, you must apply the custom policy to your group of SQL Server objects.

For this scenario, you create a static object group that contains your SQL Server virtual machines. In your own environment, you might need to create a dynamic object group so that vRealize Operations Manager discovers new SQL Server instances that become available to analyze and report on.

Prerequisites

You configured the custom policy settings for your virtual SQL Server machines. See Configure the Custom Policy Settings to Analyze and Report on vSphere Objects.

Procedure

1. To create a group type for your virtual SQL Servers, in the menu, click **Administration**, and then in the left pane and click **Configuration > Group Types**.
2. Click the **Add Group Type** icon to add a new object group type, and type **`vSphere Production Virtual Machines`**.

   You use this group type to categorize your SQL Server virtual machines for analysis.
3. In the menu, click **Environment**, and then click **Groups and Applications > Custom Groups**.

   A folder that corresponds to the group type that you just created appears in the list.
4. Click the folder named **vSphere Production Virtual Machines**, and click the plus sign to add a new object group.
5. In the New Group dialog box, add your SQL Server virtual machines.
   a. In the **Name** text box, type **`vSphere Production SQL Server Virtual Machines`**.
   b. From the **Group Type** drop-down menu, select **vSphere Production Virtual Machines**.
   c. From the **Policy** drop-down menu, select **vSphere Production Virtual SQL Servers**.
   d. In the object type drop-down menu in the Define Membership Criteria pane, expand **vCenter Adapter** and click **Virtual Machine**.
6. Click **OK** to save your object group.

   After vRealize Operations Manager collects data, the **Groups** tab displays the status for the health, risk, and efficiency of the virtual machines in the object group.

Results

You created an object type and object group to have vRealize Operations Manager analyze and report on the status of your SQL Server virtual machines.

What to do next

Create a custom dashboard so that you can view the status of your virtual SQL Servers and address problems that occur. See Dashboards.

Configure a modeling project that includes capacity planning scenarios for your production virtual SQL Servers to have vRealize Operations Manager monitor the capacity trends on these objects and notify you 60 days before your virtual SQL Servers experience capacity problems. See Chapter 6 Planning the Capacity for Your Managed Environment Using vRealize Operations Manager.

Have vRealize Operations Manager report on the CPU use and memory use of your virtual machines on a regular schedule, and send the reports to you.

Default Policy in vRealize Operations Manager

The default policy is a set of rules that applies to the majority of your objects.

The Default policy appears on the **Active Policies** tab, and is marked with the letter D in the Priority column. The Default policy can apply to any number of objects.

The Default policy always appears at the bottom in the list of policies, even if that policy is not associated with an object group. When an object group does not have a policy applied, vRealize Operations Manager associates the Default policy with that group.

A policy can inherit the Default policy settings, and those settings can apply to various objects under several conditions.

The policy that is set to Default always takes the lowest priority. If you attempt to set two policies as the Default policy, the first policy that you set to Default is initially set to the lowest priority. When you set the second policy to Default, that policy then takes the lowest priority, and the earlier policy that you set to Default is set to the second lowest priority.

You can use the Default policy as the base policy to create your own custom policy. You modify the default policy settings to create a policy that meets your analysis and monitoring needs. When you start with the Default policy, your new policy inherits all of the settings from the Default base policy. You can then customize your new policy and override these settings.

The data adapters and solutions installed in vRealize Operations Manager provide a collective group of base settings that apply to all objects. In the policy navigation tree on the **Policy Library** tab, these settings are called Base Settings. The Default policy inherits all of the base settings by default.

Policies Provided with vRealize Operations Manager

vRealize Operations Manager includes sets of policies that you can use to monitor your environment, or as the starting point to create your own policies.

Verify that you are familiar with the policies provided with vRealize Operations Manager so that you can use them in your own environment, and to include settings in new policies that you create.

Where You Find the Policies Provided with vRealize Operations Manager Policies

In the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. To see the policies provided with vRealize Operations Manager, expand the Base Settings policy.

Policies That vRealize Operations Manager Includes

All policies exist under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects. In the policy navigation tree on the **Policy Library** tab, these settings are called Base Settings.

The Base Settings policy is the umbrella policy for all other policies, and appears at the top of the policy list in the policy library. All of the other policies reside under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects.

The Config Wizard Based Policy set includes policies provided with vRealize Operations Manager that you use for specific settings on objects to report on your objects. The Config Wizard Based Policy set includes several types of policies:

- Capacity Management policies for Network I/O and Storage I/O
- Efficiency alerts policies for infrastructure objects and virtual machines
- Health alerts policies for infrastructure objects and virtual machines
- Overcommit policies for CPU and Memory
- Risk alerts policies for infrastructure objects and virtual machines

The Default Policy includes a set of rules that applies to the majority of your objects.

The VMware Management Policies set includes policies that you use for your type of environment, such as production as opposed to test and development. These policies contain settings that monitor for peak periods, batch and interactive workloads, and demand and allocation models. The VMware Management Policies set provided with vRealize Operations Manager include the following policies:

Table 4-77. Functions of VMware Management Policies

VMware Management Policy	What it does
VMware Excludes over-sized analysis	Does not calculate reclaimable capacity from oversized virtual machines
VMware Optimized for 15-minute peak periods	Configured to cause capacity alerts for workloads that spike for 15 minutes.
VMware Optimized for 30-minute peak periods	Configured to cause capacity alerts for workloads that spike for 30 minutes.
VMware Policy for Batch workloads	Optimized for batch workloads that run less than four hours.
VMware Policy for Interactive workloads	Configured to be sensitive toward interactive workloads, such as a desktop or Web server, based on 15-minute peaks with large buffers.
VMware Production Policy (Demand only)	Optimized for production loads, without using allocation limits, to obtain the most capacity.
VMware Production Policy (with Allocation)	Optimized for production loads that require the demand and allocation capacity models.
VMware Production Policy (without Allocation)	Optimized for production loads that require demand capacity models, and provides the highest overcommit without contention.
VMware Test and Dev Policy (without Allocation).	Optimized for Dev and Test environments to maximize capacity without causing significant contention, because it does not include capacity planning at the virtual machine level.

Using the Monitoring Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the monitoring policy workspace to create local policies quickly, and update the settings in existing policies. Select a base policy to use as the source for your local policy settings, and modify the thresholds and settings used for analysis and collection of data from groups of objects in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated object groups.

Prerequisites

Verify that objects groups exist for vRealize Operations Manager to analyze and collect data, and if they do not exist, create them. See Managing Custom Object Groups in VMware vRealize Operations Manager.

Procedure

1. In the menu, click **Administration**, and then in the left pane click **Policies**.
2. Click **Policy Library**, and click the **Add New Policy** icon to add a policy, or select the policy and click the **Edit Selected Policy** icon to edit an existing policy.

   You can add and edit policies on the **Policy Library** tab, and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.
3. In the Getting Started workspace, assign a name and description to the policy.

   Give the policy a meaningful name and description so that all users know the purpose of the policy.
4. Click **Select Base Policies**, and in the workspace, select one or more policies to use as a baseline to define the settings for your new local policy.

   When you create a policy, you can use any of the policies provided with vRealize Operations Manager as a baseline source for your new policy settings.
5. Click **Override Analysis Settings**, and in the workspace, filter the object types to customize your policy for the objects to associate with this policy.

   Filter the object types, and modify the settings for those object types so that vRealize Operations Manager collects and displays the data that you expect in the dashboards and views.
6. Click **Override Attributes**, and in the workspace, select the metric, property, or super metric attributes to include in your policy.

   vRealize Operations Manager collects data from the objects in your environment based on the metric, property, or super metric attributes that you include in the policy.

7 Click **Override Alert / Symptom Definitions**, and in the workspace, enable or disable the alert definitions and symptom definitions for your policy.

   vRealize Operations Manager identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

8 Click **Apply Policy to Groups**, and in the workspace, select one or more groups to which the policy applies.

   VMware vRealize Operations Manager monitors the objects according to the settings in the policy that is applied to the object group, triggers alerts when thresholds are violated, and reports the results in the dashboards, views, and reports. If you do not assign a policy to one or more object groups, VMware vRealize Operations Manager does not assign the settings in that policy to any objects, and the policy is not active. For an object group that dos not have a policy assigned, VMware vRealize Operations Manager associates the object group with the Default Policy.

9 Click **Save** to retain the settings defined for your local policy.

What to do next

After vRealize Operations Manager analyzes and collects data from the objects in your environment, review the data in the dashboards and views. If the data is not what you expected, edit your local policy to customize and override the settings until the dashboards display the data that you need.

Policy Workspace in vRealize Operations Manager

The policy workspace allows you to quickly create and modify policies. To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

How the Policy Workspace Works

Every policy includes a set of packages, and uses the defined problems, symptoms, metrics, and properties in those packages to apply to specific object groups in your environment. You can view details for the settings inherited from the base policy, and display specific settings for certain object types. You can override the settings of other policies, and include additional policy settings to apply to object types. For example, a critical production policy includes settings to track use, available resources and the time remaining on them, resource demands on the object group that determine how much stress is applied, and reclaimable capacity amounts for CPU, disk I/O, and network I/O.

Use the **Add** and **Edit** options to create policies and edit existing policies.

Where You Create and Modify a Policy

To create and modify policies, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. The policy workspace is where you select the base policies, and customize and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can apply the policy to object groups.

To remove a policy from the list, select the policy and click the red X.

Policy Workspace Options

The policy workspace includes a step-by-step workflow to create and edit a policy, and apply the policy to custom object groups.

- Getting Started Details

  When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

- Select Base Policy Details

  You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a new policy. In the policy content area, you can view the packages and elements for the base policy and additional policies that you selected to override the settings, and compare the differences in settings highlighted between these policies. You select the settings and objects types to display.

- Analysis Settings Details

  You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

- Workload Automation Details

  You can set the workload automation options for your policy, so that vRealize Operations Manager can balance the workload in your environment per your definition.

- Collect Metrics and Properties Details

  You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

- Alert and Symptom Definitions Details

  You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

- Custom Profiles Details

  Custom profiles show how many more of a specified object can fit in your environment depending on the available capacity and object configuration. You can enable or disable custom profiles for your policy.

- Apply Policy to Groups Details

  You can assign your local policy to one or more groups of objects to have VMware vRealize Operations Manager analyze those objects according to the settings in your policy, trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

Getting Started Details

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

Where You Assign the Policy Name and Description

To add a name and description to a policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Getting Started**. The name and description appear in the workspace.

Table 4-78. Name and Description Options in the Add or Edit Monitoring Policy Workspace

Option	Description
Name	Name of the policy as it appears in the Add or Edit Monitoring Policy wizard, and in areas where the policy applies to objects, such as Custom Groups.
Description	Meaningful description of the policy. For example, use the description to indicate which policy is inherited, and any specific information that users need to understand the relationship of the policy to one or more groups of objects.
Start with	The base policy that will be used as a starting point. All settings from the base policy will be inherited as default settings in your new policy. You can override these settings to customize the new policy. Select a base policy to inherit the base policy settings as a starting point for your new policy.

Select Base Policy Details

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a new policy. In the policy content area, you can view the packages and elements for the base policy and additional policies that you selected to override the settings, and compare the differences in settings highlighted between these policies. You select the settings and objects types to display.

How the Select Base Policies Workspace Works

To create a policy, select a base policy from which your new custom policy inherits settings. To override some of the settings in the base policy according to the requirements for the service level agreement for your environment, you can select and apply a separate policy for a management pack solution. The override policy includes specific settings defined for the types of objects to override, either manually or that an adapter provides when it is integrated with vRealize Operations Manager. The settings in the override policy overwrite the settings in the base policy that you selected.

When you select and apply a policy in the left pane to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the applied policy history list in the right pane.

The right pane displays tabs for the inherited policy configuration, and your policy, and displays a preview of the selected policy tab in the Policy Preview pane. When you select one of the policy tabs, you can view the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of enabled and disabled changes.

In the right pane, you select the objects to view so that you can see which policy elements apply to the object type. For example, when you select the StorageArray object type, and you click the tab to display the configuration settings for your policy, the Policy Preview pane displays the local packages for the policy and the object group types with the number of policy elements in each group.

You can preview the policy settings for all object types, only the object types that have settings changed locally, or settings for new object types that you add to the list, such as Storage Array storage devices.

Where You Select and Override Base Policies Settings

To select a base policy to use as a starting point for your own policy, and to select a policy to override one or more settings that your policy inherits from the base policy, in the menu, select **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left add a name for the policy and click **Select Base Policy**. The policy configuration, objects, and preview appear in the workspace.

Table 4-79. Base Policy and Override Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Show changes for	Select the objects to view changes. ■ All object types. Displays the number of enabled and disabled alert definitions, symptom definitions, and metrics and properties, the number of enabled and disabled changes, and the object type groups and the number of local policy elements for each group. ■ All object types with overrides. Displays the object types that have changes applied, with the objects types selected for override. Use the drop-down menu to select object types. Click the filter button to add the selected object type to the list so that you can preview and configure the settings. ■ Add settings for new set of objects. Provides a list of the object types so that you can select an object type, such as **Storage Devices > SAN**, and add the selected object to the Object types list.
Override settings from additional policies	Select and apply one or more policies to override the settings that your policy inherits from the base policy.
Apply	Applies the override policy to your policy, and lists the override policy in the applied policy history.
Applied policy template history	Displays the policies that you selected to override the settings in your policy.
Configuration inherited from base policy	When selected, displays a preview of the inherited policy configuration in the Policy Preview pane.
Configuration settings defined in this policy	When selected, displays a preview of your policy configuration in the Policy Preview pane.
Policy Preview	Displays summary information about the local packages and object group types. ■ Packages (Local). Displays the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of policy elements for each object group. ■ Object Type groups. Displays the associated object groups. ■ Drop down arrows on packages and settings. Displays the packages and settings for the displayed policies.

Analysis Settings Details

You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

How the Analysis Settings Workspace Works

When you turn on and configure the analysis settings for a policy, you can override the settings for the policy elements that vRealize Operations Manager uses to trigger alerts and display data. These types of settings include badge score symptom thresholds based on alerts, situational settings such as committed projects to calculate capacity and time remaining, and other detailed settings.

You expand a policy element setting and configure the values to make your policy specific. For example, to reclaim capacity, you can set percentages to have vRealize Operations Manager indicate when a resource is oversized, idle, or powered off.

Policies focus on objects and object groups. When you configure policy element settings for your local policy, you must consider the object type and the results that you expect to see in the dashboards and views. If you do not make any changes to the settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

Where You Set the Policy Analysis Settings

To set the analysis settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Analysis Settings**. The analysis settings for host systems, virtual machines, and other object types that you select appear in the workspace.

Table 4-80. Analysis Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Show changes for	Select the objects to view changes. ■ All object types. Displays the number of enabled and disabled alert definitions, symptom definitions, and metrics and properties, the number of enabled and disabled changes, and the object type groups and the number of local policy elements for each group. ■ All object types with overrides. Displays the object types that have changes applied, with the objects types selected for override. Use the drop-down menu to select object types. Click the filter button to add the selected object type to the list so that you can preview and configure the settings. ■ Add settings for new set of objects. Provides a list of the object types so that you can select an object type, such as **Storage Devices > SAN**, and add the selected object to the Object types list.
Right pane - Analysis Settings for object types	The right pane displays a list of the object types that you selected in the left pane. Expand a view of the policy elements and settings for the object type so that you can have vRealize Operations Manager analyze the object type. Expand the view for the object type so that you can view and modify the threshold settings for the following policy elements: ■ Workload ■ Anomaly ■ Fault ■ Capacity and Time Remaining ■ Stress ■ Compliance ■ Reclaimable Capacity ■ Density ■ Time Range Click the lock icon on the right of each element to override the settings and change the thresholds for your policy.

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. You can then override the settings and have vRealize Operations Manager calculate the metrics for CPU use

and memory use, and display the demand for resources on the selected objects, based on your settings.

How the Workload Element Works

The Workload element determines how vRealize Operations Manager reports on the resources that the selected object group uses. The resources available to the object group depend on the amount of configured and usable resources.

- A specific amount of physical memory is a configured resource for a host system, and a specific number of CPUs is a configured resource for a virtual machine.
- The usable resource for an object or an object group is a subset of, or equal to, the configured amount.
- The configured and usable amount of a resource can vary depending on the type of resource and the amount of virtualization overhead required, such as the memory that an ESX host machine requires to run the host system. When accounting for overhead, the resources required for overhead are not considered to be usable, because of the reservations required for virtual machines or for the high availability buffer.

Where You Override the Policy Workload Element

To view and override the policy workload analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The workload settings for the object types that you selected appear in the right pane.

View the Workload policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-81. Policy Workload Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Workload drop-down menu	When expanded, displays a list of the resource containers. You can enable or disable the resource containers for the workload calculation.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Workload** to view the badge score symptom thresholds for Workload policy settings for a selected object, as defined in the policy applied to the object.

Table 4-82. Policy Workload Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Workload drop-down menu	When expanded, displays a list of the resource containers. You can enable or disable the resource containers for the workload calculation.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then in the left pane select an object. Click **Analysis > Workload** to view the badge score symptom thresholds for Workload policy settings for a selected object, as defined in the policy applied to the object.

Policy Anomaly Element

An anomaly is an unusual or abnormal event that occurs on an object. You can turn on and configure the settings for the Anomaly element for the object types in your policy so that you can override the settings and have vRealize Operations Manager determine the acceptable level of abnormal behavior for an object according to the historical metrics data for that object, based on your settings.

Where You Override the Policy Anomaly Element

To view and override the policy Anomaly analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The anomalies settings for the object types that you selected appear in the right pane.

View the Anomaly 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 4-83. Policy Anomaly Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Anomalies** to view the badge score symptom thresholds for Anomalies policy settings, as defined in the policy applied to the object.

Policy Fault Element

A fault is an object-based error condition, such as `Guest file system out of space` for a virtual machine, or `Host connectivity` for a host system. You can turn on and configure the settings for

the Fault element for the object types in your policy so that you can override the settings and have vRealize Operations Manager determine and quantify the severity of problems experienced by selected objects, based on your settings.

Where You Override the Policy Fault Element

To view and override the policy Fault analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The fault settings for the object types that you selected appear in the right pane.

View the Fault 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 4-84. Policy Fault Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Overrides button	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment** and then from the left pane select an object. Click **Analysis > Faults** to view the badge score symptom thresholds for Faults policy settings for a selected object, as defined in the policy applied to the object.

Policy Capacity Remaining and Time Remaining Element

Capacity is a measurement of the amount of memory, CPU, and disk space for an object. Time remaining is a measure of the amount of time left before your objects run out of capacity. You can turn on and configure the settings for the Capacity Remaining and Time Remaining element for the object types in your policy so that you can override the settings and have vRealize Operations Manager report on the available capacity remaining and the amount of time remaining before resources run out, based on your settings.

How the Capacity Remaining and Time Remaining Element Works

The Capacity Remaining and Time Remaining element determines how vRealize Operations Manager reports on the available capacity and time until resources run out for a specific object type group.

- The capacity remaining indicates the capability of your environment to accommodate new machines. vRealize Operations Manager calculates the capacity remaining as a percentage of the overall capacity that remains for the number of virtual machines, compared to the total number of virtual machines that can be deployed on the selected object.

- The time remaining indicates the amount of time that remains before the object group consumes all of the resources. vRealize Operations Manager calculates the time remaining as the number of days remaining until all of the capacity would be consumed, minus the number of days allocated to the provisioning buffer.
- Usable capacity is a measurement of the percentage of capacity available, minus the capacity affected when you use high availability, and you set capacity buffer amounts on the memory, CPU, network, datastore, and disk space buffers. If you set overcommit values, the usable capacity measurement adds the capacity to the amount of usable capacity available.
- You can modify the usable capacity settings to use high availability so that vRealize Operations Manager provides enough objects and resources to address throughput and any potential loss of data. You can also modify the calculation type and the buffer rules.
- Capacity settings for resource containers are enabled or disabled for analysis. For the Memory, CPU, and Disk Space resource containers, you can enable or disable the demand and allocation. For the Network I/O resource container, you can enable or disable the data transmit rate, data receive rate, and the use rate. For the Datastore I/O resource container, you can enable or disable the outstanding I/O requests, reads and writes per second, and the read and write rate. You can also enable or disable the vSphere configuration limit.
- The peak consideration setting causes vRealize Operations Manager to apply stress settings to account for peak uses in capacity.
- You can have vRealize Operations Manager account for committed projects that you defined so that you can plan the future capacity of your objects. Because committed projects are scenarios that forecast the future capacity of objects, accounting for committed projects affects the time remaining score.
- The number of days set for the provisioning time buffer is based on the amount of time you require to provision the objects in your environment, from the time of ordering those objects to deploying them. To keep the Time Remaining score above zero, your objects must have more days of capacity available than the provisioning time buffer.

Where You Override the Policy Capacity Remaining and Time Remaining Element

To view and override the policy Capacity Remaining and Time Remaining analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The capacity remaining and time remaining settings for the object types that you selected in the workspace appear in the right pane.

View the Capacity Remaining and 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 4-85. Policy Capacity Remaining and Time Remaining Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Symptom Thresholds for Time Remaining Score and Capacity Score	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. The badge score symptom threshold for the Capacity Remaining and Time Remaining policy settings appear on the following tabs for a selected object: ■ ■ In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Capacity Remaining** ■ In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Time Remaining**

Table 4-85. Policy Capacity Remaining and Time Remaining Element Settings in the Add or Edit Monitoring Policy Workspace (continued)

Option	Description
Usable Capacity settings for resource containers	Displays the selected resource containers and resources to include in the analysis, overcommit types and values for resources such as memory and CPU, and the capacity buffer percentage for each resource container. ■ Capacity Buffer %. Defines the percentage of capacity reserved on virtual machines so that the virtual machine does not consume all of its resources. The capacity buffers are defined on Cluster objects and Host objects to reserve some resources for failover. ■ Overcommit. Displays the over commitment type, such as memory or CPU. ■ Value. Displays the amount of over commitment on capacity resources. To change these settings, select a resource container and double-click the value you want to change.
Additional settings that affect time and capacity remaining calcilations	The available settings depend on the object type you select. ■ ■ High Availability. When selected, vRealize Operations Manager report on the capacity available to the object type group. You can have vRealize Operations Manager take into consideration the High Availability (HA) settings. ■ Peak Consideration. When selected, vRealize Operations Manager includes the Stress element in the capacity remaining and time remaining calculations. ■ Committed Projects. When selected, if you committed one or more projects on an object type, and added capacity scenarios to those projects to plan for future capacity requirements, vRealize Operations Manager accounts for the committed projects in the capacity remaining and time remaining calculations. ■ Capacity Calculation. Indicates on what status vRealize Operations Manager reports. You can select either the current value or a trend of values as the basis for the capacity analysis. ■ Provisioning Time Buffer. Indicates the number of days allowed to provision physical or virtual resources. vRealize Operations Manager uses this number to calculate the capacity remaining and time remaining for resource types, and shortens the time remaining scores. Peak consideration, committed projects, and the provisioning buffer settings, as defined in the applied policy, appear on the following tabs for a selected object. ■ In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Capacity Remaining**. ■ In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Time Remaining**.

Policy Stress Element

Stress is a measurement of the workload on an object over time, including CPU, memory, network I/O, and datastore I/O. You can turn on and configure the settings for the Stress element for the object types in your policy so that you can override the settings and have vRealize Operations Manager analyze the resources used for an object or object group over a period of time, and report the historical workload based on your settings.

How the Stress Element Works

The Stress element determines how vRealize Operations Manager reports on the demand for resources and usable capacity over time.

- When you include the Stress element in your policy, you can use the stress score to identify hosts and machines that require additional resources, and identify hosts that require fewer virtual machines, so that you can avoid performance problems in your environment.
- When you select Peak Consideration in the Capacity & Time Remaining element, vRealize Operations Manager can use the Stress element to account for peaks in capacity usage.
- Stress is the percentage of demand over time, where stress extends above the stress noise line. For example, the stress line might be 70 percent of the percentage of workload over time, based on the setting used for exceeding the demand. As vRealize Operations Manager calculates capacity and time remaining, you might want to account for these spikes and peaks.

  To set the stress settings, use the sliding analysis settings. The settings for stress might need to differ between policies used to monitor your infrastructure versus virtual machines. For example, for an infrastructure policy, your recommended levels for the stress settings might be 10 (warning), 30 (immediate), and 50 (critical). For virtual machines, the settings might be 5 (warning), 10 (immediate), and 20 (critical). For a test and development policy, you might want vRealize Operations Manager to trigger an alert when the level reaches 10 percent. For a production policy, you typically want to ensure that sufficient capacity exists for peak use.

Where You Override the Policy Stress Element

To view and override the policy Stress analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The stress settings for the object types that you selected appear in the right pane.

View the Stress 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 4-86. Policy Stress Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Symptom Thresholds for Time Remaining Score and Capacity Score	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Stress** to view the badge score symptom thresholds for Stress policy settings for a selected object, as defined in the policy applied to the object.
Stress settings for resource containers	Displays the resource container and settings for exceeding the demand during the time range defined in the policy Time element. In the menu, click **Environment**, and then from the left pane select an object. Click **Analysis > Stress** to view the percentage for exceeding demand for a selected object, as defined in the applied policy. The Sliding Analysis Window defines the time period that vRealize Operations Manager checks for stress, which occurs at the defined range of minutes, or during the entire range defined for the data range in the Time policy element, to monitor for peak stress periods. To modify the setting, select the resource container setting, such as **Disk Space > Usage**, double-click the Sliding Analysis Window setting, and select either **Any** or **Entire Range**. When the setting is Any, you can modify the Minute Peak value to an interval in minutes for vRealize Operations Manager to monitor your objects and report on peak times of stress.

Policy Compliance Element

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the Compliance element for the object types in your policy. You can override the base policy settings and have vRealize Operations Manager report on the compliance results for virtual machines and related objects, such as the ratios of virtual machines to hosts, memory demand, and CPU demand.

Where You Override the Policy Compliance Element

To view and override the policy Compliance analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The compliance settings for the object types that you selected appear in the right pane.

View the Compliance policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-87. Policy Compliance Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then in the left pane select and an object. Click **Analysis > Compliance** to view the badge score symptom thresholds for Compliance policy settings for a selected object, as defined in the policy applied to the object.

Policy Reclaimable Capacity Element

Reclaimable capacity is a measurement of the CPU, memory, and disk space for your objects that is designated as waste. You can turn on and configure the settings for the Reclaimable Capacity element for the object types in your policy so that you can override the settings and have vRealize Operations Manager analyze and report on the capacity that you can reclaim from unused or underused objects. You can then provision the reclaimed capacity to other objects in your environment, based on your settings.

How the Reclaimable Capacity Element Works

The Reclaimable Capacity element determines how vRealize Operations Manager reports the amount of reclaimable capacity of objects such as CPU, memory, and disk space for each object in your environment.

When you include the reclaimable capacity element in your policy, you can use the reclaimable capacity score to identify the amount of resources that can be reclaimed and provisioned to other objects.

Where You Override the Policy Reclaimable Capacity Element

To view and override the policy Reclaimable Capacity analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The reclaimable capacity settings for the object types that you selected appear in the right pane.

View the Reclaimable Capacity 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 4-88. Policy Reclaimable Capacity Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then in the left pane select an object. Click **Analysis > Reclaimable Capacity** to view the badge score symptom thresholds for Reclaimable Capacity policy settings for a selected object, as defined in the policy applied to the object.
Reclaimable capacity settings for resource containers	Displays the configurable percentages for vRealize Operations Manager to use to report when a resource is determined to be oversized, idle, or powered off. In the menu, click **Environment**, and then in the left pane select an object. Click **Analysis > Reclaimable Capacity** to view the settings for disk and CPU idle levels, and the percentages used to determine when a resource is oversized, idle, or powered off for a selected object, as defined in the policy applied to the object. For the selected object, you can set the Oversized, Idle, and Powered Off, and Unused capacity settings. ■ An object is considered to be oversized when its recommended capacity is less than the defined percentage of its current capacity. For example, when the Oversized setting for a virtual machine is 50%, the virtual machine is considered to be oversized when its capacity is half of the current capacity available. ■ An object is considered to be idle when the object operates below the idle level for the defined percentage of time. For example, when the CPU idle level is set to 100 MHz for a virtual machine, and the flag for the idle level is set to 90%, the virtual machine is considered to be idle when the speed of its CPU drops below 100 MHz for 90% of the time. ■ An object is flagged as powered off when the object is powered down for the defined percentage of time. For example, when the powered off flag is set to 90%, a virtual machine is flagged as powered off when it is powered down at least 90% of the time. ■ An object is considered to be unused when its timestamp attribute has not changed for the defined number of days, which means that the object has not been accessed. For example, when the flag for the disk space reclaimable snapshot space is set to 60 days for a virtual machine, if the virtual machine or the files on it have not been accessed for 60 days, the virtual machine is considered to be unused.

Policy Density Element

Density is a measurement of the sizing ratio of your objects based on available CPU as opposed to demand, and available memory as opposed to demand. You can unlock and configure the settings for the Density element for the object types in your policy. You can override the base policy settings and have vRealize Operations Manager report on the density results for virtual machines and related objects, such as the ratios of virtual machines to hosts, memory demand, and CPU demand. For example, to reduce the virtual machine density on a host machine, move some of the virtual machines to another host.

Where You Override the Policy Density Element

To view and override the policy Density analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The density settings for the object types that you selected appear in the right pane.

View the Density 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 4-89. Policy Density Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Badge Score Symptom Threshold	Set the symptom threshold values for the policy element to levels that update the badge scores to meet the criteria for your environment. vRealize Operations Manager uses the symptom threshold values to trigger alerts that appear in the Alerts Overview and Dashboard scores. In the menu, click **Environment**, and then in the left pane select an object. Click **Analysis > Density** to view the badge score symptom thresholds for Density policy settings for a selected object, as defined in the policy applied to the object.

Policy Time Element

Time indicates the schedule and range of days and hours that vRealize Operations Manager monitors the use of resources for your objects, and the maintenance schedule selected for periodic and repeatable maintenance. You can turn on and configure the settings for the Time element for the object types in your policy so that you can override the settings and have vRealize Operations Manager report on metrics and calculate analytics for the group at specific times.

How the Time Element Works

The Time element determines when and how vRealize Operations Manager tracks resources on a specific object type.

Where You Override the Policy Time Element

To view and override the policy Time analysis setting, in the menu, click **Administration**, and then in the left pane, click **Policies**. Click the **Policy Library** tab. Click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, click **Analysis Settings**, then select one or more objects in the left pane. The time settings for the object types that you selected appear in the right pane.

View the Time 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 4-90. Policy Time Element Settings in the Add or Edit Monitoring Policy Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Track Usage	Determines the times when vRealize Operations Manager runs the capacity analytics calculations. ■ At all times. Monitor the amount of time tracked 24 hours a day, 7 days a week. ■ Specific days and times. Select when to track the time use.
Data Range	Sets the number of days to include in the analysis of time use.
Maintenance Schedule	Sets a time to perform maintenance tasks. During maintenance times, vRealize Operations Manager does not calculate analytics.

Workload Automation Details

You can set the workload automation options for your policy, so that vRealize Operations Manager can balance the workload in your environment per your definition.

How the Workload Automation Workspace Works

You click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option your policy inherits the parent policy settings. The graphic on the right updates to reflect your changes.

Where You Set the Policy Workload Automation

To set the workload automation for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Workload Automation**.

Table 4-91. Workload Automation in the Add or Edit Monitoring Policy Workspace

Option	Description
Balance Workloads	Select how vRealize Operations Manager balances the workload. Select Aggressive balancing when you have stable populations. It minimizes contention but moves workloads more, which can cause disruption. Select Conservative balancing when you have dynamic populations. It exposes potential contention, but moves workloads less.
Consolidate Workloads	Select how vRealize Operations Manager combines the workload. The consolidation policy setting does not affect the placement of virtual machines across clusters. ■ Select more consolidation when you have populations with steady demand. It puts workloads into as few hosts as possible to reduce licensing and power costs. However, this approach might cause less responsive capacity. ■ Select less consolidation when you have population with irregular demand. It uses all available hosts, which leaves more room for demand spikes. However, this approach increases licensing and power costs.

Table 4-91. Workload Automation in the Add or Edit Monitoring Policy Workspace (continued)

Option	Description
Cluster Headroom	Headroom establishes a required capacity buffer, for example, 20 percent, that can reduce the possibility of clusters running out of CPU capacity or memory, but can also cause a rebalance to occur sooner than it might have otherwise.
Advanced Settings	Click **Advanced Settings** to select what type of virtual machines vRealize Operations Manager moves first to address workload. This option also allows you to set Storage vMotion on or off. The default is ON.

Collect Metrics and Properties Details

You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

How the Collect Metrics and Properties Workspace Works

When you create or customize a policy, you can override the base policy settings to have vRealize Operations Manager collect the data that you intend to use to generate alerts, and report the results in the dashboard scores.

To define the metric and super metric symptoms, metric event symptoms, and property symptoms, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Symptom Definitions**.

Where You Override the Policy Attributes

To override the attributes and properties settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Collect Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

Table 4-92. Collect Metrics and Properties Options

Option	Description
Actions	Select one or more attributes and select enable, disable, or inherit to change the state and KPI for this policy.
Filter options	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.
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 Definitions Details

You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

How the Alert and Symptom Definitions Workspace Works

vRealize Operations Manager collects data for objects and compares the collected data to the alert definitions and symptom definitions defined for that object type. Alert definitions include associated symptom definitions, which identify conditions on attributes, properties, metrics, and events.

You can configure your local policy to inherit alert definitions from the base policies that you select, or you can override the alert definitions and symptom definitions for your local policy.

Before you add or override the alert definitions and symptom definitions for a policy, familiarize yourself on the available alerts and symptoms.

- To view the available alert definitions, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Alert Definitions**.
- To view the available symptom definitions, in the menu, click **Alerts** and then in the left pane click **Alert Settings > Symptom Definitions**. Symptom definitions are available for metrics, properties, messages, faults, smart early warnings, and external events.

A summary of the number of problem and symptoms that are enabled and disabled, and the difference in changes of the problem and symptoms as compared to the base policy, appear in the Analysis Settings pane of the policies workspace.

Where You Override the Alert Definitions and Symptom Definitions

To override the alert definitions and symptom definitions for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Alert / Symptom Definitions**. The definitions appear in the workspace.

Policy Alert Definitions and Symptom Definitions

You can override the alert definitions and symptom definitions for each policy.

- Policy Alert Definitions

  Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

- Policy Symptom Definitions

  Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

How the Policy Alert Definitions Work

vRealize Operations Manager uses problems to trigger alerts. A problem manifests when a set of symptoms exists for an object, and requires you to take action on the problem. Alerts indicate problems in your environment. vRealize Operations Manager generates alerts when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert occurs, vRealize Operations Manager presents the triggering symptoms for you to take action.

Some of the alert definitions include predefined symptoms. When you include symptoms in an alert definition, and enable the alert, an alert is generated when the symptoms are true.

The Alert Definitions pane displays the name of the alert, the number of symptoms defined, the adapter, object types such as host or cluster, and whether the alert is enabled as indicated by **Local**, disabled as indicated by **not Local**, or inherited. Alerts are inherited with a green checkmark by default, which means that they are enabled.

You can automate an alert definition in a policy when the highest priority recommendation for the alert has an associated action.

To view a specific set of alerts, you can select the badge type, criticality type, and the state of the alert to filter the view. For example, you can set the policy to send fault alerts for virtual machines.

Where You Modify the Policy Alert Definitions

To modify the alerts associated with policies, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Alert / Symptom Definitions**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-93. Alert Definitions in the Add or Edit Monitoring Policy Workspace

Option	Description
Actions	Select one or more alert definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	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 4-93. Alert Definitions in the Add or Edit Monitoring Policy 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. ■ Name. Meaningful name for the alert definition. ■ Symptom Definitions. Number of symptoms defined for the alert. ■ Actionable Recommendations. Only recommendations with actions in the first priority, as they are the only ones you can automate. ■ Automate. When the action is set to Local, the action is enabled for automation when an alert triggers. Actions that are enabled for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to **Local** with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ State. Alert definition state, either enabled as indicated by **Local**, disabled as indicated by **not Local**, or inherited from the base policy.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

How the Policy Symptom Definitions Work

vRealize Operations Manager uses symptoms that are enabled to generate alerts. When the symptoms used in an alert definition are true, and the alert is enabled, an alert is generated.

When a symptom exists for an object, the problem exists and requires that you take action to solve it. When an alert occurs, vRealize Operations Manager presents the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

To assess objects for symptoms, you can include symptoms packages in your policy for metrics and super metrics, properties, message events, and faults. You can enable or disable the symptoms to determine the criteria that the policy uses to assess and evaluate the data collected from the objects to which the policy applies. You can also override the threshold, criticality, wait cycles, and cancel cycles.

The Symptoms pane displays the name of the symptom, the associated management pack adapter, object type, metric or property type, a definition of the trigger such as for CPU usage, the state of the symptom, and the trigger condition. To view a specific set of symptoms in the package, you can select the adapter type, object type, metric or property type, and the state of the symptom.

When a symptom is required by an alert, the state of the symptom is enabled, but is dimmed so that you cannot modify it. The state of a required symptom includes an information icon that you can hover over to identify the alert that required this symptom.

Where You Modify the Policy Symptom Definitions

To modify the policy package of symptoms, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left, click **Alert / Symptom Definitions**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-94. Symptom Definitions in the Add or Edit Monitoring Policy Workspace

Option	Description
Actions	Select one or more symptom definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	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 4-94. Symptom Definitions in the Add or Edit Monitoring Policy 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. ■ Name. Symptom definition name as defined in the list of symptom definitions in the Content area. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ Type. Object type on which the symptom definition must be evaluated. ■ Trigger. Static or dynamic threshold, based on the number of symptom definitions, the object type and metrics selected, the numeric value assigned to the symptom definition, the criticality of the symptom, and the number of wait and cancel cycles applied to the symptom definition. ■ State. Symptom definition state, either enabled, disabled, or inherited from the base policy. ■ Condition. Enables action on the threshold. When set to Override, you can change the threshold. Otherwise set to default. ■ Threshold. To change the threshold, you must set the State to **Enabled**, set the condition to **Override**, and set the new threshold in the Override Symptom Definition Threshold dialog box.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Custom Profiles Details

Custom profiles show how many more of a specified object can fit in your environment depending on the available capacity and object configuration. You can enable or disable custom profiles for your policy.

Where You Set the Policy Custom Profiles

To apply the policy to object groups, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Custom Profiles**.

Table 4-95. Custom Profiles Options

Option	Description
Actions	Select one or more profiles and select enable, disable, or inherit to change the state for this policy.
Filter options	Deselect the options in the **State** drop-down menu, to narrow the list of attributes. ■ Enabled. Indicates that a profile will be calculated. ■ Enabled (Force). Indicates state change due to a dependency. ■ Disabled. Indicates that a profile not be calculated. ■ Inherited. Indicates that the state of this profile is inherited from the base policy and will be calculated. ■ Inherited. Indicates that the state of this profile is inherited from the base policy and will not be calculated.
Object Type	Filters the profiles list by object type.

Apply Policy to Groups Details

You can assign your local policy to one or more groups of objects to have VMware vRealize Operations Manager analyze those objects according to the settings in your policy, trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

How the Apply Policy to Groups Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more object groups. VMware vRealize Operations Manager uses the settings in the policy to analyze and collect data from the associated objects, and displays the data in dashboards, views, and reports.

Where You Apply a Policy to Groups

To apply the policy to object groups, in the menu, click **Administration**, and then in the left pane click **Policies**. Click the **Policy Library** tab, and click the **Add New Policy** icon to add a policy or click the **Edit Selected Policy** icon to edit a policy. In the Add or Edit Monitoring policy workspace, on the left click **Apply Policy to Groups**.

Apply Policy to Groups Options

To apply the policy to groups of objects, select the check box for the object group in the workspace.

You can then view the details about each object group associated with the policy. In the menu, click **Administration**, and then in the left pane click **Policies**. Click **Active Policies > Related Objects**. Click an object group in the list of groups, and view the summary in the Details pane.

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**, 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 4-96. 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 with the exception of 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.

Table 4-96. Define Monitoring Goals Questions (continued)

Option	Description
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

You can set compliance on your objects to meet the defined standards and determine the compliance of your objects against the configuration standards.

Defining Compliance Standards

Compliance is used to monitor the vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches in your environment to ensure that the settings on your objects meet the defined standards.

vRealize Operations Manager includes alerts for *VMware vSphere Hardening Guide* versions 6.0 and 5.5. vRealize Operations Manager generates compliance alerts when symptoms trigger on your vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches.

To enforce compliance on virtual machines, vRealize Operations Manager includes several compliance risk profiles. You apply the risk profiles to groups of virtual machines based on whether you must ensure a high, medium, or low level of security in your environment.

- Risk Profile 1 includes all available compliance rules as symptoms, and enforces the highest level of security for your virtual machines. This profile is enabled by default.
- Risk Profile 2 enforces a medium level of security for your environment, and includes fewer symptoms than Risk Profile 1. This profile is disabled by default.
- Risk Profile 3 enforces a low level of security, and includes fewer symptoms than Risk Profile 2. This profile is disabled by default.

All the compliance standards in vRealize Operations Manager, including any standards that you define, are based on alert definitions. The generated alerts and symptoms appear as violations to the compliance standards on the **Analysis > Compliance** tab for a selected object.

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

The following video is an example of how you can now ensure compliance of your VMware vSphere 6.0 and 5.5 objects, including your vCenter Server instances, ESXi hosts, virtual machines, distributed port groups, and distributed virtual switches. The compliance alerts include definitions and symptoms, and are based on the compliance rules in the vSphere Hardening Guides 6.0 and 5.5.

vRealize Operations Manager 6.3 Compliance for vSphere 6.0 Objects (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrom6.3_compliance_vsphere6_objects)

vRealize Operations Manager Compliance for vSphere 6.0 Objects

To ensure compliance of your vSphere 6.0 and 5.5 objects, vRealize Operations Manager includes compliance alerts for *VMware vSphere Hardening Guide* versions 6.0 and 5.5. These hardening guide alerts are now based on object type.

When you customize a policy to enable the *vSphere Hardening Guide* alerts, you can enable vSphere 6.0 and 5.5 alerts for the following object types and versions:

- ESXi host is violating *vSphere Hardening Guide* (5.5 and 6.0)
- vCenter Server is violating *vSphere Hardening Guide* (6.0)
- Virtual machine is violating Risk Profile 1 in *vSphere Hardening Guide* (5.5 and 6.0)
- Virtual machine is violating Risk Profile 2 in *vSphere Hardening Guide* (5.5 and 6.0)
- Virtual machine is violating Risk Profile 3 in *vSphere Hardening Guide* (5.5 and 6.0)
- vSphere Distributed Port Group is violating *vSphere Hardening Guide* (6.0)
- vSphere Distributed Virtual Switch is violating *vSphere Hardening Guide* (6.0)

By default, the alert named `Virtual machine is violating Risk Profile 1` is the only active alert among the risk profiles. You can configure this profile later, and choose one of the other risk profiles.

To determine whether an alert triggered against *vSphere Hardening Guide* 6.0 or 5.5, you must examine the underlying symptoms. For example, for the alert named `ESXi Host is violating vSphere Hardening Guide`, the following underlying symptoms for the alert include:

- ESXi.set-account-lockout - The count failed login attempts before the account is locked out exceeded maximum (*vSphere Hardening Guide* 6.0)
- DCUI service is running (*vSphere Hardening Guide* 5.5)

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

Reset Default Content to Ensure Current Compliance Standards for vSphere 6.0 and 5.5 Objects

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 the option to overwrite alert definitions and symptom definitions.

If you do not overwrite your alert definitions and symptom definitions with the new content provided with this release, some compliance rules will include the new alert and symptom definitions, while other compliance rules will continue to use outdated alert and symptom definitions.

User Scenario: Ensure Compliance of Your vSphere 6.0 Objects

As the virtual infrastructure administrator for your company, you must ensure that your vSphere 6.0 objects comply with the compliance rules in the *vSphere Hardening Guide*. You use the compliance alerts in vRealize Operations Manager to monitor your objects for violations to your compliance standards. When a compliance alert triggers on your vCenter Server instance, hosts, virtual machines, distributed port groups, or distributed switches, you investigate the compliance violation. You must and resolve the violation so that the violated object continues to meet industry security standards.

You manage and monitor the security of your production, test, and development environments. Your objects consist of multiple vCenter Server instances, with hosts, virtual machines, distributed port groups, and distributed switches in each instance.

Your CIO requires that you run SSH on all vCenter Server instances and host machines in your production and test environments. You monitor all hosts to ensure that they comply with the SSH requirement. You produce a compliance report each week to prove to your manager and the compliance team that your objects comply with the implemented security standards.

To enforce and report on the compliance of your vSphere 6.0 objects, you enable the compliance rules in the *vSphere Hardening Guide*. Then, you enable the appropriate alerts, and apply a risk profile to your virtual machines. After vRealize Operations Manager collects the compliance data from your objects, you resolve any rule violations that occurred, and create a report of the compliance results for your manager and the compliance team.

The Alert definitions provided with vRealize Operations Manager are based on object types instead of the specific versions of the hardening guides. To use these alerts, you no longer must create a custom group and apply the policy to that group.

Some alert definitions are common between vSphere 6.0 and vSphere 5.5 objects. vRealize Operations Manager checks vSphere 6.0 symptoms against 6.0 objects, 5.5 symptoms against 5.5 objects, and a combination of 6.0 and 5.5 symptoms against both versions of the objects.

Prerequisites

Verify that the current version of vRealize Operations Manager is installed and running.

Procedure

1. In vRealize Operations Manager, enable the compliance rules.
   a. Click **Administration**, and click **Solutions**.
   b. Click the VMware vSphere solution, and click **Configure**.
   c. In the Manage Solution dialog box, click **Define Monitoring Goals**.

d Under **Enable vSphere Hardening Guide Alerts**, click **Yes** and click **Save**.

e When vRealize Operations Manager reports that the default policy is configured to collect compliance data on your objects, click **OK** and click **Close**.

2 Enable the compliance alert definitions in the default policy.

a In the menu, click **Administration**, and then on the left pane click **Policies** and then click the **Policy Library** tab.

b Click the **Default Policy**, and click **Edit Selected Policy**.

c In the Edit Monitoring Policy workspace on the left, click **Alert / Symptom Definitions**.

d In the filter text box in the Alert Definitions pane, enter **`hardening`**.

Several alert definitions appear, which you use to enforce compliance on your objects. Each alert displays the number of symptoms and the object type to which the alert applies. You can see the alert definitions for risk profiles 1, 2, and 3, which you use to ensure high, medium, or low security on your virtual machines.

e Click the alert named `vCenter is violating vSphere Hardening Guide`.

f In the State column, click the down arrow, and select **Local**.

g To enable compliance alerts on your virtual machines, distributed port groups, and distributed switches, enable the other alert definitions, and click **Save**.

3 View the symptom set in the alert definition for the ESXi host.

a In the menu, click **Alerts** and then in the left pane, click **Alert Definitions**.

b In the filter text box, enter **`hardening`**.

c Click the alert named `vCenter is violating vSphere Hardening Guide`.

d In the lower pane, locate the alert impact, criticality, and symptom set.

e Scroll through the symptom set and examine the symptoms, which can trigger an alert, for the host.

f Below the symptom set, examine the recommendation to fix the problem if this alert triggers on your host.

g Click the link to the *VMware vSphere Hardening Guide*.

The Web page opens to the list of *VMware vSphere Security Hardening Guides* at http://www.vmware.com/security/hardening-guides.html.

4 Focus in on the alerts for the host in your production vCenter Server instance.

a In the menu, click **Alerts** and then in the left pane, click **All Alerts**.

b Click the link in the compliance alert named `ESXi Host is violating vSphere Hardening Guide`.

c Examine the symptoms, which displays the hosts that violated the rules in the *vSphere Hardening Guide*.

d For the first host listed, click on the host name, and examine the violations on the Summary tab.

5 Run a report for your compliance team.

a In the menu, click **Administration**, and then in the left pane click **Reports**.

b In the filter text box, enter **`hardening`**.

The report named *VMware vSphere Hardening Guide - Non-compliance Report* appears.

c On the Report Templates tab, click **Run Template**, and wait for vRealize Operations Manager to generate the report.

d Click **Generated Reports**.

The report appears, and provides PDF and CSV versions for you to download.

e In the Download column, click the **PDF** icon and examine the content in the report.

The non-compliance report appears for the host, and includes the date and time that you ran the report. It also identifies you as the user who ran the report. The report displays the noncompliant rules that ran on the object and its descendants. In the report, you can see the criticality and status of the alert, the object name, and the type on which the alert triggered.

f In the Download column, click the **CSV** icon, and examine the content of the spreadsheet.

The spreadsheet provides an easy way to see a summary of the results, and allows you to import the data into another application.

Results

You have ensured that the compliance rules, are enforced on the objects in your vCenter Server instances, according to the *VMware vSphere Hardening Guide*.

What to do next

To examine the compliance alert definitions for your other objects, click **Alerts > Alert Definitions**.

User Scenario: Define a Compliance Standard for Custom Standards

As a virtual infrastructure administrator, you are responsible for the vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches in your environment. To ensure the compliance of your vSphere objects, you create a compliance standard based on an alert definition.

In vRealize Operations Manager, you can configure an alert definition to use as a compliance standard. Any alert definition that you configure with the subtype named Compliance appears on the **Compliance** tab. For more information on Compliance Tab, see Compliance Tab.

When you create an alert definition as a compliance standard, you add all the relevant symptom definitions to the alert definition. Each symptom is a rule in the compliance standards. For most alert definitions, you must avoid adding too many symptoms to the alert definition.

vRealize Operations Manager includes alerts for *VMware vSphere Hardening Guide* versions 6.0 and 5.5.

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

In this scenario, the alert notifies you when SSH is not running on the host.

Procedure

1 Configure Basic Information for the Host Compliance Standard

To create an alert definition that is also a compliance standard, you first configure the name, base object type, and the alert impact.

2 Add Symptoms to the Host Compliance Standard

You add symptoms and recommendations to the alert definition so that when the host system compliance alert is generated, the symptoms appear as rules on the Compliance tab.

Configure Basic Information for the Host Compliance Standard

To create an alert definition that is also a compliance standard, you first configure the name, base object type, and the alert impact.

The name of the alert is the name of the standard on the Compliance tab.

Procedure

1 In the menu, click **Alerts** and select **Alert Definitions** from the **Alert Settings** drop-down menu in the left pane.

2 Click the plus sign to add a definition.

3 Type a name and description.

In this scenario, enter **`Organization Host Compliance Standards`**.

4 Click **Base Object Type**, expand **vCenter Adapter** in the drop-down menu, and select **Host System**.

5 Click **Alert Impact** and configure the metadata for this alert definition.

a From the **Impact** drop-down menu, select **Risk**.

b From the **Criticality** drop-down menu, select **Symptom Based**.

c From the **Alert Type and Subtype** drop-down menu, expand **Virtualization/Hypervisor** and select **Compliance**.

Any alert where you use the Compliance subtype is processed as a compliance standard.

d Configure the **Wait Cycle** and **Cancel Cycle** with a value of **1**.

What to do next

Add the symptoms that act as the compliance rules. See Add Symptoms to the Host Compliance Standard.

Add Symptoms to the Host Compliance Standard

You add symptoms and recommendations to the alert definition so that when the host system compliance alert is generated, the symptoms appear as rules on the Compliance tab.

Prerequisites

Configure the name, host object type, and alert impact setting for the alert so that it appears as a compliance standard. See Configure Basic Information for the Host Compliance Standard.

Procedure

1. In the **Alert Definition Workspace** window, click **Add Symptom Definitions** and add the SSH symptom.
   a. From the **Symptom Definition Type** drop-down menu, select **Metric / Property**.
   b. In the **Symptom** search text box, enter **SSH**.
   c. Drag the symptom named SSH `service is running` to the symptoms workspace.

   If you add multiple symptoms for your own scenario, and you determine that the alert must trigger when any of the symptoms occur, you would select **Any** from the drop-down menu named **This symptom set is true when**.
2. In the workspace navigation pane, click **Add Recommendations**, and create a recommendation for the standard.
   a. Click the plus sign to add a recommendation.
   b. Enter a name for the recommendation in the text box.

      For example, enter **`Turn on the SSH service`**. If you have a local runbook, you can provide a link to your local instructions.
   c. Click **Save**.
   d. Drag the recommendation to the workspace.

   In your own scenario, you can create multiple recommendations for the standard.
3. Click **Save**.

Results

If the symptom condition becomes true, the symptom is triggered and the compliance alert is generated for the object. Because the alert definition includes the subtype named Compliance, the generated alert appears as a compliance standard on the Compliance tab.

What to do next

Review the Compliance tab for standards that indicate that other objects are out of compliance, including vCenter Server instances, virtual machines, distributed port groups, and distributed switches. See Compliance Tab.

Configuring Super Metrics

The super metric is a mathematical formula that contains one or more metrics. It is a custom metric that you design to help track combinations of metrics, 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 4-97. 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 4-97. 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 System Audit for vRealize Operations Manager.
- 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.

Create a Super Metric

Create a super metric when you want to check the health of your environment, but cannot find a suitable metric to perform the analysis.

Procedure

1. On the menu, click **Administration** and in the left pane click **Configuration > Super Metrics**.
2. Click the **Add** icon.
3. Enter a meaningful name for the super metric such as **SM–AvgVMCPUUsage%** in the **Name** text box.

4 Define the formula for the super metric.

Select the function or operator. Select the metrics or attribute types that you want to use in each function or with each operator. For example, to add a super metric that captures average CPU usage across all virtual machines, perform the following tasks.

a For Function, select **avg**.

b In the **Operators** text box, select the left parenthesis, then select the right parenthesis. Click between the two parentheses to position the cursor.

c Optionally, click the **This object** icon.

If the **This object** icon is selected, the super metric is calculated using the metrics and properties of the object you use for the super metric. If the icon is not selected, the super metric is calculated using metrics and properties of other objects and the super metric function displays the object with a long description.

d In the **Adapter Type** text box of the Object Types pane, select **vCenter Adapter**.

e From the list of object types, select **Virtual Machine**

f In the **Metrics** pane, expand the CPU category, scroll down, and double-click the **Usage (%)** metric.

The formula appears as a mathematical function. To view the formula in a textual format, click the **Show Formula Description** icon. If the formula syntax is wrong, an error message appears. The formula ends with `depth=1`. With `depth=1`, you assign the super metric to an object type that is one level above virtual machines in the relationship chain. The super metric appears as a metric for that object type. With `depth=2`, you assign the super metric to an object type that is two levels above virtual machines, for example a Cluster.

5 To assign the super metric to an object type at `depth=1`, type 2 instead of 1, so that `depth=2` is displayed.

6 Verify that the super metric formula has been created correctly.

a Click the **Visualize Super Metric** icon.

b In the Objects pane, double-click one of the objects listed.

A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

7 Click **Save**.

8 Associate the super metric with an object. vRealize Operations Manager calculates the super metric for the target objects and displays it as a metric for the object type.

a In the Super Metrics workspace, select the super metric.

b In the **Object Types** tab, click the **Add** icon.

c In the Select Object Type text box, select the required object. For example, if you created your super metric for Host Systems under the vCenter Adapter, expand **vCenter Adapter**, and select **Host Systems**.

d Click **Select**.

After one collection cycle, the super metric appears on each of the objects of the specified object type. For example, you defined the super metric to calculate average CPU usage across all virtual machines and assigned it to the Host System object type. After one collection cycle, the super metric appears as a super metric on each host.

What to do next

In the **Policies > Edit Policy > Attributes** workspace, you must select and enable each super metric. See Custom Policies. Wait at least one collection cycle for the super metric to begin collecting and processing data. Then review your super metric in the **All Metrics** tab.

Enhancing Your Super Metrics

You can enhance your super metrics by using clauses and resource entry aliasing.

Where Clause

The **where** clause verifies whether a particular metric value can be used in the super metric. Use this clause to point to a different metric of the same object, such as **where = "metric_group|my_metric > 0**.

For example:

```
count(${adaptertype = ExampleAdapter, objecttype = ExampleObject, metric =
ExampleGroup|Rating, depth=2, where = "==1"})
```

Resource Entry Aliasing

Resource entries are used to retrieve metric data from vRealize Operations Manager for computing super metrics. A resource entry is the part of an expression which begins with **$** followed by a **{..} block**. When computing a super metric, you might have to use the same resource entry multiple times. If you have to change your computation, you must change every resource entry, which might lead to errors. You can use resource entry aliasing to rewrite the expression.

The following example, shows a resource entry that has been used twice.

```
(min(${adapterkind=VMWARE, resourcekind=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=”>=0”}) + 0.0001)/(max(${adapterkind=VMWARE,
resourcekind=HostSystem, attribute=cpu|demand|active_longterm_load, depth=5,
where=”>=0”}) + 0.0001)"
```

The following example shows how to write the expressing using resource entry aliasing. The output of both expressions is the same.

```
(min(${adapterkind=VMWARE, resourcekind=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=">=0"} as cpuload) + 0.0001)/(max(cpuload) +
0.0001)"
```

Follow these guidelines when you use resource entry aliasing:

- When you create an alias, make sure that after the resource entry you write **as** and then **alias:name**. For example: **${...} as alias_name**.
- The alias cannot contain the ()[]+-*/%|&!=<>,.?:$ special characters, and cannot begin with a digit.
- An alias name, like all names in super metric expressions, is case-insensitive.
- Use of an alias name is optional. You can define the alias, and not use it in an expression.
- Each alias name can be used only once. For example: `${resource1,…} as r1 + ${resource2,…} as R1`.
- You can specify multiple aliases for the same resource entry. For example: **${...} as a1 as a2**.

Conditional Expression ?: Ternary Operators

You can use a ternary operator in an expression to run conditional expressions.

For example: `expression_condition ? expression_if_true : expression_if_false`.

The result of the conditional expression is converted to a number. If the value is not 0, then the condition is assumed as true.

For example: `–0.7 ? 10 : 20` equals 10. `2 + 2 / 2 – 3 ? 4 + 5 / 6 : 7 + 8` equals 15 (7 + 8).

Depending on the condition, either **expression_if_true** or **expression_if_false** is run, but not both of them. In this way, you can write expressions such as, `${this, metric=cpu|demandmhz} as a != 0 ? 1/a : –1`. A ternary operator can contain other operators in all its expressions, including other ternary operators.

For example: `!1 ? 2 ? 3 : 4 : 5` equals 5.

Exporting and Importing a Super Metric

You can export a super metric from one vRealize Operations Manager instance and import it to another vRealize Operations Manager instance. For example, after developing a super metric in a test environment, you can export it from the test environment and import it use in a production environment.

If the super metric to import contains a reference to an object that does not exist in the target instance, the import fails. vRealize Operations Manager returns a brief error message and writes detailed information to the log file.

Procedure

1 Export a super metric.

   a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.

   b Select the super metric to export, click the **Actions** icon and select **Export Selected Super Metric** icon.

   vRealize Operations Manager creates a super metric file, for example, `SuperMetric.json`.

   c Download the super metric file to your computer.

2 Import a super metric.

   a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.

   b Click the **Actions** icon and select **Import Super Metric**.

   c (Optional). If the target instance has a super metric with the same name as the super metric you are importing, you can either overwrite the existing super metric or skip the import, which is the default.

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 4-98. 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 Collect 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 4-99. 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 4-100. 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 4-101. 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(${adapterkind=*adaptkind*, resourcekind=*reskind*, resourcename=*resname*, identifiers={id1=*val1*id2=*val2*,...}, metric=*a*|*b:instance*|*c*})	Returns a single data point of a particular metric for the *resname* specified in the argument. This super metric does not take values from the children or parents of the object.
funct(${adapterkind=*adaptkind*, resourcekind=*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 4-101. Looping Function Formats (continued)

Argument syntax example	Description
funct(${adapterkind=*adaptkind*, resourcekind=*reskind*, depth=*dep*}, metric=*a*|*b:optional_instance*|*c*})	Returns a set of data points. This function iterates metrics of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects, where *depth* describes the object location in the relationship chain. For example, a typical relationship chain includes a data center, cluster, host, and virtual machines. The data center is at the top and the virtual machines at the bottom. If the super metric is assigned to the cluster and the function definition includes depth = 2, the super metric takes values from the virtual machines. If the function definition includes depth = -1, the super metric takes values from the data center.
funct(${adapterkind=*adaptkind*, resourcekind=*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(${adapterkind=VMWARE, resourcekind=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 4-102. 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.
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.

Table 4-102. Single Functions (continued)

Function	Format	Description
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 to enclose or insert between functions.

Table 4-103. 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.
( )	Parentheses
[ ]	Use in an array of expressions
[x, y, z]	An array containing x, y, z. For example, min([x, y, z])

Configuring Objects

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.

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 check if a datastore is connected or providing data, or you can power on a virtual machine.

Object Discovery

Its ability to monitor and collect data on objects in your systems environment makes vRealize Operations Manager a critical tool in maintaining system uptime and ensuring ongoing good health for all system resources from virtual machines to applications to storage - across physical, virtual and cloud infrastructures.

Following are examples of objects that can be monitored.

- vCenter Server
- Virtual machines
- Servers/hosts
- Compute resources
- Resource pools
- Data centers
- Storage components
- Switches
- Port groups
- Datastores

Adapters – Key to Object Discovery

vRealize Operations Manager collects data and metrics from objects using adapters, the central components of management packs, which in turn make up vRealize Operations Manager solutions. When you configure the vSphere Solution, for example, you create adapter instances customized for your environment with unique names, port numbers, and so on. You must create an adapter instance for each vCenter Server in your deployment.

Locate existing adapters in the UI as follows: in the menu, click **Administration**, then click **Solutions** in the left pane.

As shown in the screenshot, the Solutions screen lists available solutions at the top of the screen. When you select a solution, the available adapters appear in the lower half of the screen. Existing adapter instances related to each adapter are listed in the second column.

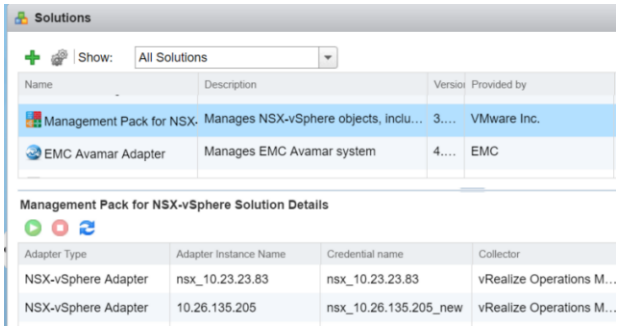

For complete information on configuring management packs and adapters, see Connecting vRealize Operations Manager to Data Sources

When you create a new adapter instance, it begins discovering and collecting data from the objects designated by the adapter, and notes the relationships between them. Now you can begin to manage your objects.

About Objects

Objects are the structural components of your mission-critical IT applications: virtual machines, datastores, virtual switches and port groups are examples of objects.

Because downtime equals cost - in unused resources and lost business opportunities - it's crucial that you successfully identify, monitor and track objects in your environment. The goal is to proactively isolate, troubleshoot and correct problems even before users are aware that anything is wrong.

When a user actually reports an issue, the solution should be quick and comprehensive.

For a complete list of objects that can be defined in vRealize Operations Manager refer to Object Discovery.

vRealize Operations Manager gives you visibility into objects including applications, storage and networks across physical, virtual and cloud infrastructures through a single interface that relates performance information to positive or negative events in the environment.

Managing Objects

When you monitor a large infrastructure, the number of objects and corresponding metrics in vRealize Operations Manager grows rapidly, especially as you add solutions that extend dynamic monitoring and alerts to more parts of your infrastructure. vRealize Operations Manager gives you ample tools to stay abreast of events and issues.

Adding Objects and Configuring Object Relationships

vRealize Operations Manager automatically discovers objects and their relationships once you create an adapter instance. You have the added ability to manually add any objects that you want monitored and to configure object relationships using abstract concepts rather than the connections recorded by vRealize Operations Manager. Where vRealize Operations Manager might discover the classic parent-child relationships between objects, you can create relationships between objects that might not normally be related. For example, you could configure all the datastores supporting a company department to be related.

When objects are related, a problem with one object appears as an anomaly on related objects. So object relationships can help you to identify problems in your environment quickly. The object relationships that you create are called custom groups.

Custom Groups

To create an automated management system you need some way to organize objects so that you can quickly gain insights. You can achieve a high level of automation using custom groups. You have multiple options for tailoring group attributes to support your monitoring strategy.

For example, you can designate a group either to be static or to be updated automatically with membership criteria that you designate. Consider a non-static group of all virtual machines that are powered on and have OS type Linux. When you power on a new Linux VM, it is automatically added to the group and the policy is applied.

For additional flexibility, you can also specify individual objects to be always included or excluded from a given custom group. Or you can have a different set of alerts and capacity calculations for your production environment versus your testing environments.

Managing Applications

vRealize Operations Manager allows you to create containers or objects that can contain a group of virtual machines or other objects in different structural tiers. This new application can then be managed as a single object, and have health badges and alarms aggregated from the child objects of the group.

For example, the system administrator of an online training system might request that you monitor components in the Web, application and database tiers of the training environment. You build an application that groups related training objects together in each tier. If a problem occurs with one of the objects, it is highlighted in the application display and you can investigate the source of the problem

The Power of Object Management

Using the power of object management, including metrics and alerts - some prepackaged into dashboards and policies, others that you combine into custom monitoring tools - you'll keep a close watch on the objects, applications and systems that must stay up and running.

Managing Objects in Your Environment

An object is the individual managed item in your environment for which vRealize Operations Manager collects data, such as a router, switch, database, virtual machine, host, and vCenter Server instances.

The system requires specific information about each object. When you configure an adapter instance, vRealize Operations Manager performs object discovery to start collecting data from the objects with which the adapter communicates.

An object can be a single entity, such as a database, or a container that holds other objects. For example, if you have multiple Web servers, you can define a single object for each Web server and define a separate container object to hold all of the Web server objects. Groups and applications are types of containers.

Categorize your objects using tags, so that you can easily find, group, or filter them later. A tag type can have multiple tag values. You or vRealize Operations Manager assigns objects to tag values. When you select a tag value, vRealize Operations Manager displays the objects associated with that tag. For example, if a tag type is Lifecycle and tag values are Development, Test, Pre-production, and Production, you might assign virtual machine objects VM1, VM2, or VM3 in your environment to one or more of these tag values, depending on the virtual machine function.

Adding an Object to Your Environment

You might want to add an object by providing its information to vRealize Operations Manager. For example, some solutions cannot discover all the objects that might be monitored. For these solutions, you must either use manual discovery or manually add the object.

When you add an individual object, you provide specific information about it, including the kind of adapter to use to make the connection and the connection method. For example, a vSAN adapter does not know the location of the vSAN devices that you want to monitor.

Prerequisites

Verify that an adapter is present for the object you plan to add. See Installing Optional Solutions in vRealize Operations Manager

Procedure

1. In the menu, click **Administration**, then select **Configuration > Inventory Explorer** from the left pane.
2. On the toolbar, click the plus sign.
3. Use the topic menus to reveal all fields and provide the required information.

Option	Description
Display name	Enter a name for the object. For example, enter **vSAN–Host1**.
Description	Enter any description. For example, enter **vSAN–Host monitored with vSAN adapter**

Option	Description
Adapter type	Select an adapter type. For example, select **vSAN Adapter**.
Adapter instance	Select an adapter instance.
Object type	Select an object type. For a vSAN adapter, you might select vSAN-Host. When you select the object type, the dialog box selections change to include information you provide so that vRealize Operations Manager can find and connect with the selected object type.
Host IP address	Enter the host IP. For example, enter the IP address of vSAN-Host1.
Port number	Accept the default port number or enter a new value.
Credential	Select the Credential, or click the plus sign to add new login credentials for the object.
Collection interval	Enter the collection interval, in minutes. For example, if you expect the host to generate performance data every 5 minutes, set the collection interval to 5 minutes.
Dynamic Thresholding.	Accept the default, Yes.

4 Click **OK** to add the object.

Results

vSAN-Host1 appears in the Inventory Explorer as a host object type for the vSAN adapter type.

What to do next

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 Explorer: 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.

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 anomaly on related objects. So object relationships can help you to identify problems in your environment quickly.

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 4-104. 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 Explorer** in the left pane.
2. Click the **Manage Tags** icon above the list of tags.
3. Click the **Add New Tag** icon to add a new row and type the name of the tag in the row.

   For example, type **`Cloud Objects`** and click **Update**.
4. With the new tag selected, click the **Add New Tag Value** icon to add a new row and type the name of the value in the row.

   For example, type **`Video Cloud`** and click **Update**.
5. Click **OK** to add the tag.
6. Click the tag to which you want to add objects to display the list of object tag values.

   For example, click **Cloud Objects** to display the Video Cloud object tag value.
7. Drag objects from the list in the right pane of the Inventory Explorer onto the tag value name.

   You can press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.

   For example, if you want to assign datacenters that are connected through the vCenter Adapter, type **`vCenter`** in the search filter and select the datacenter objects to add.

Use a Tag to Find an Object

The quickest way to find an object in vRealize Operations Manager is to use tags. Using tags is more efficient than searching through the entire object list.

Tag values that can also be tags are Applications and Object Types. For example, the Object Types tag has values for each object that is in vRealize Operations Manager, such as Virtual Machine, which includes all the virtual machine objects in your environment. Each of these virtual machines is also a tag value for the Virtual Machine tag. You can expand the tag value list to select the value for which you want to see objects.

Procedure

1. In the menu, click **Administration**, then click **Configuration > Inventory Explorer** in the left pane.
2. In the tag list in the center pane, click a tag for an object with an assigned value.

   When you click a tag, the list of values expands under the tag. The number of objects that is associated with each value appears next to the tag value.

A plus sign next to a tag value indicates that the value is also a tag and that it contains other tag values. You can click the plus sign to see the subvalues.

3 Select the tag value.

The objects that have that tag value appear in the pane on the right. If you select multiple tag values, the objects in the list depend on the values that you select.

Tag Value Selection	Objects Displayed
More than one value for the same tag	The list includes objects that have either value. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, the list shows objects that have either value.
Values for two or more different tags	The list includes only objects that have all of the selected values. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, and you also select an adapter instance such as vC-1 of the vCenter Adapter instance tag, only Datacenter or Host System objects associated with vC-1 appear in the list. Datacenter or Host System objects associated with other adapter instances do not appear in the list, nor do objects that are not Datacenter or Host System objects.

4 Select the object from the list.

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 Explorer list.

Where You Find Manage Object Tags

In the menu, click **Administration**, then click **Configuration > Inventory Explorer** 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 Explorer: 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 clilck **Configuration > Inventory Explorer** 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 Explorer: 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 **Configuration > Inventory Explorer**. The system lists all the objects in your environment.

Inventory Explorer List Options

The center pane includes object tag options. The right pane includes toolbar options for all of the objects in your environment.

Table 4-105. 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 4-105. 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 4-106. Inventory Explorer 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 4-106. Inventory Explorer 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. See Summary Tab.
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 Explorer** 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 4-107. 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 4-107. 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 Explorer** 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 Explorer** 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 Discover Objects Workspace.

Table 4-108. 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 4-109. 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 4-109. 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 Explorer** in the left pane. Click **Start Maintenance** in the List tool bar.

Table 4-110. 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.

Inventory Explorer: 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 click **Configuration > Inventory Explorer** 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 VMware vRealize Operations Manager

A custom object group is a container that includes one or more objects. vRealize Operations Manager uses custom groups to collect data from the objects in the group, and report on the data collected.

Why Use Custom Object Groups?

You use groups to categorize your objects and have the system collect data from the groups of objects and display the results in dashboards and views according to the way you define the data to appear.

You can create static groups of objects, or dynamic groups with criteria that determine group membership as vRealize Operations Manager discovers and collects data from new objects added to the environment.

vRealize Operations Manager provides commonly used object group types, such as World, Environment, and Licensing. The system uses the object group types to categorize groups of objects. You assign a group type to each group so that you can categorize and organize the groups of objects that you create.

Types of Custom Object Groups

When you create custom groups, you can use rules to apply dynamic membership of objects to the group, or you can manually add the objects to the group. When you add an adapter, the groups associated with the adapter become available in vRealize Operations Manager.

- Dynamic group membership. To dynamically update the membership of objects in a group, define rules when you create a group. vRealize Operations Manager adds objects to the group based on the criteria that you define.
- Mixed membership, which includes dynamic and manual.
- Manual group membership. From the inventory of objects, you select objects to add as members to the group.

- Groups associated with adapters. Each adapter manages the membership of the group. For example, the vCenter Server adapter adds groups such as datastore, host, and network, for the container objects in the vSphere inventory. To modify these groups, you must do so in the adapter.

Administrators of vRealize Operations Manager can set advanced permissions on custom groups. Users who have privileges to create groups can create custom groups of objects and have vRealize Operations Manager apply a policy to each group to collect data from the objects and report the results in dashboards and views.

When you create a custom group, and assign a policy to the group, the system uses the criteria defined in the applied policy to collect data from and analyze the objects in the group. vRealize Operations Manager reports on the status, problems, and recommendations for those objects based on the settings in the policy.

Note Only custom groups defined explicitly by users can be exported from or imported to vRealize Operations Manager. Users are able to export or import multiple custom groups. Once an import function has been executed, the user must check to determine if a policy or policies should be associated with the imported group. Export-import operations are available for user defined (created explicitly by user) custom groups only.

How Policies Help vRealize Operations Manager Report On Object Groups

vRealize Operations Manager analyzes the objects in the object group and reports on the workload, capacity, stress, anomalies, and faults of the object group, among other attributes.

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 plus sign 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.

The Groups data grid displays an overview of the state of each group.

Table 4-111. 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.

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, and 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 4-112. 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 4-112. 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 white list. ■ 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.
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 blacklist. ■ 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.

Table 4-112. New Group Workspace (continued)

Option	Description
	■ 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.

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 or all 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.

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 plus sign to add an application.

The Applications data grid displays an overview of the state of each application.

Table 4-113. 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 plus sign.
3. Click **Basic n-tier Web App** and click **OK**.

   The Application Management page that appears has two rows. Select objects from the bottom row to populate the tiers in the top row.
4. Type a meaningful name such as **`Online Training Application`** in the Application text box.

5 For each of the Web, application and database tiers listed, add the objects to the Tier Objects section.

   a Select a tier name. This is the tier that you populate.

   b To the left of the object row, select object tags to filter for objects that have that tag value. Click the tag name once to select the tag from the list and click the tag name again to deselect the tag from the list. If you select multiple tags, objects displayed depend on the values that you select.

      You can also search for the object by name.

   c To the right of the object row, select the objects to add to the tier.

   d Drag the objects to the Tier Objects section.

6 Click Save to save the application.

Results

The new application appears in the list of applications on the Environment Overview Applications page. If any of the components in any of the tiers develops a problem, the application displays a yellow or red status.

What to do next

To investigate the source of the problem, click the application name and see Evaluating Object Information Using Badge Alerts and the Summary Tab.

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 4-114. 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 4-115. 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

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.

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 generally 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 4-116. 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.py file import reskndmetric YourCustomFilename.xml`

The file is imported into global storage and is accessible from the supported widgets.

4 If you update an exisiting file and must re-import the file, append `--force` to the above import command and run it.

For example, `./vcops-cli.py file import reskndmetric YourCustomFilename.xml --force`.

What to do next

To verify that the XML file is imported, configure one of the supported widgets and ensure that the new file appears in the drop-down menu.

You can also create a custom set of metrics to display the widgets, from the Manage Metric Configuration.

Widget Definitions List

A widget is a pane on a dashboard that contains information about configured attributes, resources, applications, or the overall processes in your environment. Widgets can provide a holistic, end-to-end view of the health of all of the objects and applications in your enterprise. If your user account has the necessary access rights, you can add and remove widgets from your dashboards.

Table 4-117. 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	Shows a chart of the Capacity values for a specific resources over the past 7 days.
Capacity Utilization	Shows the capacity or workload utilization for objects so that you can identify problems with capacity and workload. Indicates objects that are underutilized, optimal, and overutilized, and indicates why they are constrained.
Container Details	Shows the health and alert counts for each tier in a single selected container.
Container Object List	Shows a list of all defined resources and object types.
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.

Table 4-117. Summary of Widgets (continued)

Widget Name	Description
Density	Shows the density breakdown as charts for the past 7 days for a specific resource.
DRS Cluster Settings	Shows the workload of the available clusters and the associated hosts.
Efficiency	Shows the status of the efficiency-related alerts for the objects that it is configured to monitor. Efficiency is based on generated efficiency alerts in your environment.
Environment	Lists the number of resources by object or groups them by object type.
Environment Overview	Shows the performance status of objects in your virtual environment and their relationships. You can click an object to highlight its related objects and double-click an object to view its Resource Detail page.
Environment Status	Shows statistics for the overall monitored environment.
Faults	Shows a list of availability and configuration issues for a selected resource.
Forensics	Shows how often a metric had a particular value, as a percentage of all values, within a given time period. It can also compare percentages for two time periods.
Geo	Shows where your objects are located on a world map, if your configuration assigns values to the Geo Location object tag.
Health	Shows the status of the health-related alerts for the objects that it is configured to monitor. Health is based on generated health alerts in your environment.
Health Chart	Shows health information for selected resources, or all resources that have a selected tag.
Heat Map	Shows a heat map with the performance information for a selected resource.
Mashup Chart	Brings together disparate pieces of information for a resource. It shows a health chart, an anomaly count graph, 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.
Reclaimable Capacity	Shows a percentage chart representing the amount of reclaimable capacity for a specific resource that has consumers.
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.

Table 4-117. Summary of Widgets (continued)

Widget Name	Description
Scoreboard	Shows values for selected metrics, which are typically KPIs, with color coding for defined value ranges.
Scoreboard Health	Shows color-coded health or workload scores for selected resources.
Sparkline Chart	Shows graphs that contain metrics for an object . If all of the metrics in the Sparkline Chart widget are for an object that another widget provides, the object name appears at the top right of the widget.
Stress	Shows a weather map of the average stress over the past 6 weeks for a specific resource.
Tag Picker	Lists all defined resource tags.
Text Display	Reads text from a Web page or text file and shows the text in the user interface.
Time Remaining	Shows a chart of the Time Remaining values for a specific resources over the past 7 days.
Top Alerts	Lists the alerts most likely to negatively affect your environment based on the configured alert type and objects.
Top-N	Shows the top or bottom N number metrics or resources in various categories, such as the five applications that have the best or worth health score.
Topology Graph	Shows multiple levels of resources between nodes.
View	Shows a defined view depending on the configured resource.
Weather Map	Uses changing colors to show the behavior of a selected metric over time for multiple resources.
Workload	Shows workload information for a selected resource.

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.

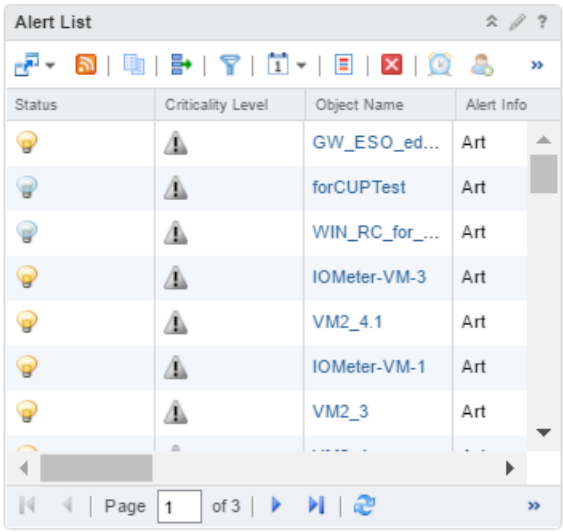

Where You Find the Alert List Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Alert List Widget and Configuration Options

The alert list widget includes toolbar options, data grid options and configuration options.

Option	Description
Dashboard Navigation	Actions you can run on the selected alert. For example, you use the option to open a vCenter Server, datacenter, 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.
RSS Feed	Send an RSS feed of the alert to your Web browser. Only alerts that appear in the widget as it is configured are included. For example, if the widget is set to show alerts only for a particular object, only alerts for that object are included in the RSS feed. The detail message of an individual alert appears in the feed's headline. Depending on the RSS client that you use, details for all anomalies related to the alert appear in the feed's body.
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.

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.
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.
Color Row by Alert Criticality	Colors the entire row based on the criticality of the alert. ■ Red. Critical alerts. ■ Yellow. Warning or intermediate alerts
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.
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.
Filter	Locate data in the widget.

The data grid provides information on which you can sort and search.

Table 4-118. Alert List Widget Data Grid

Option	Description
Status	Current state of the alert. Possible values include Active or Canceled.
Criticality Level	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**.
Object Name	Name of the object for which the alert was generated.
Alert Info	Name of the alert definition that generated the alert.
Alert Impact	Alert badge for which the alert was generated. Possible values are Health, Risk, or Efficiency.
Object Type	Object type of the object for which the alert was generated.
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
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
Duration	Current age of the alert.
Start Time	Date and time when the alert was generated.
Update Time	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.

Table 4-118. Alert List Widget Data Grid (continued)

Option	Description
Cancel Time	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.
Control State	State of user interaction with the alert. Possible values include: ■ Open. The alert is available for action. ■ Assigned. The alert is assigned to a user for action. ■ Suspended. The alert was suspended for a specified amount of time.
User Name	Name of the user who owns the alert.

The Alert List Widget provides configuration options.

Table 4-119. Alert List Widget Configuration Options

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 .
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

Table 4-119. Alert List Widget Configuration Options (continued)

Option	Description
Tag Picker	List of defined object tags, both default and custom tags, from which you can select one or more object tag values. The objects with the selected tag values applied are the basis for the widget data. If you select more than one value for the same tag, the widget includes objects that have any of the tags applied. If you use the Tag Picker to identify data, the Selected object text box remains empty.
Filter by	Limits the alerts that appear in this alert list to those that meet the selected criteria. You can configure the following filters: ■ 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. ■ User Control State. Select one or more control states to include in the list. ■ Criticality Level Range. Select one or more levels of criticality. ■ Alert Impact. Select one or more alert badges to include in the list. ■ Time Range. Select a general date range or configure a specific date range.

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.

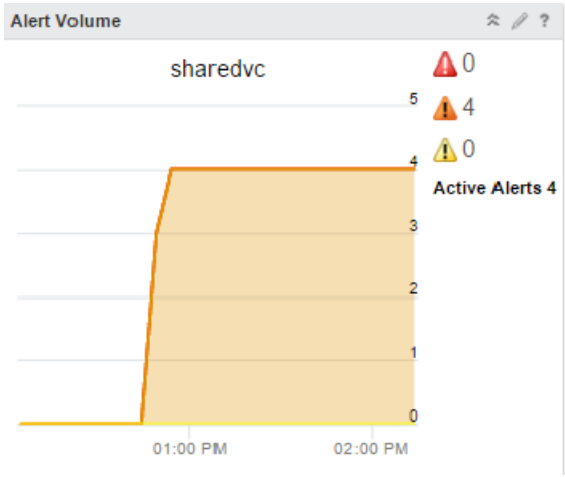

Where You Find the Alert Volume Widget and Configurations Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Table 4-120.

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.

Table 4-121.

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 .

Table 4-121. (continued)

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

Click the Anomalies score badge to go to the Anomalies analysis view the details for an anomaly.

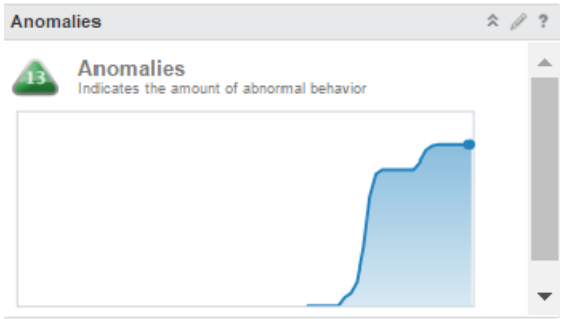

Where You Find the Anomalies Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

Table 4-122.

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	■ 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.
Selected Object	Object that is the basis for the widget data.
Object List	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. 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. If you select an object in the list, the object becomes the selected object for the widget.

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

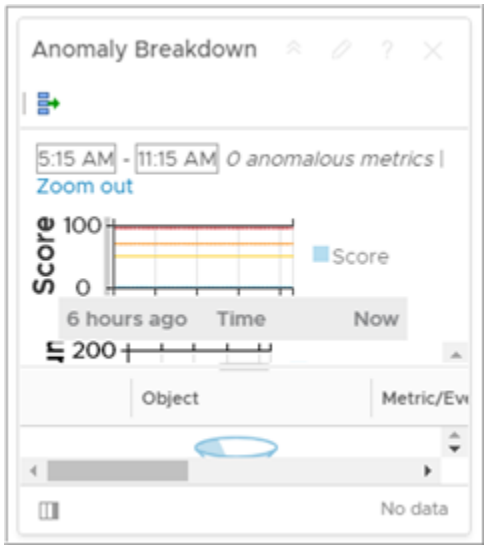

You can add the Anomaly Breakdown widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Where You Find the Anomaly Breakdown Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

Table 4-123. Anomaly Breakdown Widget Options

Option	Description
Score	Badge 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.

To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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.
Mode	Display a single or multiple objects.
Show	Select the number of objects to display when in Multiple mode.

Option	Description
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Object list	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.

Capacity Remaining Widget

The Capacity Remaining widget displays a score indicating the remaining computing resources as a percent of the total consumer capacity.

Where You Find the Capacity Remaining Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

Capacity Remaining Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Selected Object	Object that is the basis for the widget data.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

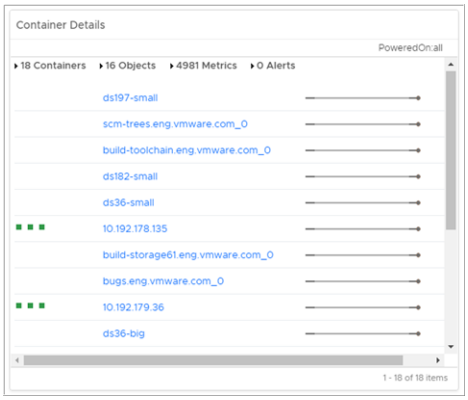

How the Container Details Widget and Configuration Options Work

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 and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

Container Details Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Mode	You can change the size of the graph using the Compact or Large buttons.
Object tree	You can filter the list of objects in the object data grid. You can select one or more object types and all objects from this type are displayed in the data grid. For example, if you want to observe information about the VMs and vCenter Server in the inventory, you can click on **Collapse All** and select **Virtual Machine** and **vCenter Server** from the object tree . As a result, the data grid shows only VMs and vCenter Server objects from the inventory.
Object data grid	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. When you select an object from the list it appears in the selected object pane. **Note** You can select to observe only one object from the inventory
Selected Object	Object that is the basis for the widget data.

Capacity Utilization Widget

The Capacity Utilization widget displays a visual summary of the capacity and workload resources used by the objects in your environment.

How the Capacity Utilization Widget and Configuration Options Work

Use the Capacity Utilization widget to identify which objects are underutilized, overutilized, and operating at optimum capacity levels.

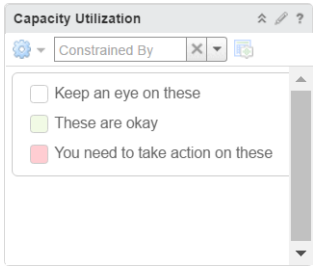

The Capacity Utilization widget appears with the name Current Object Utilization on the dashboard named Workload Utilization, which is provided with vRealize Operations Manager.

When you point to an object, vRealize Operations Manager displays a popup summary that displays the object name, the capacity used by the object, and the reason that the capacity resource is constrained on the object. To display the analysis details about the capacity on the object so that you can further troubleshoot the problem, click **Details**. By default the objects are constrained by the most constrained metric.

For example, if the capacity of your cluster is greater than 100 percent because it is constrained by disk space, click **Details** to display the **Analysis > Capacity Remaining** tab and analyze the capacity remaining for the cluster. On this tab, you can determine whether the memory or disk space being consumed on the object is causing the overutilization problem.

When several objects are affected, an object icon displays the number of objects in the utilization summary. The number of objects appears next to the utilization labels.

For example, a host object icon might display 12 to indicate the number of overutilized hosts in your environment. To display the individual affected hosts, point to the host object icon. A list of host machines appears, including the individual host names and links, the percentage of capacity used on each host, and the reason why the capacity is constrained. To further analyze the capacity details for each host, click the host link to display the **Analysis > Capacity Remaining** tab so that you can further troubleshoot the problem.

When numerous objects are affected, a graph reflects the number of objects in the utilization summary. The number of objects appears next to the utilization labels.

You can use the Capacity Utilization widget to ensure that all of your objects are as close to optimal as possible. The metric calculation displays a value that indicates the margin by which the object is away from optimal usage. The resolution depends on the object type. For a consumer object, such as a virtual machine, the resolution is usually to right-size the object to bring it to optimal. For a provider object, such a cluster, you can determine whether you must add capacity or move your existing workloads to reduce stress in the environment.

You can add the Capacity Utilization widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Where You Find the Capacity Utilization Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The Capacity Utilization configuration options appear when you click **Edit** in the Capacity Utilization widget. The Workload Utilization dashboard, which is provided with vRealize Operations Manager, displays the Capacity Utilization widget with the name Current Object Utilization.

To customize the data that appears in the dashboard widget, click **Content** in the left pane, and click **Dashboards**. On the Dashboards toolbar, click the plus sign to add a dashboard or the pencil to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the pencil to access the configuration options.

Capacity Utilization Widget and Configuration Options

The Capacity Utilization widget includes toolbar and configuration options.

Option	Description
Action	Displays the available actions for a specific object. For example, if you select the host object icon, the Action icon is enabled and displays all the available actions you can carry out. Some of the options are: **Power Off VM**, P**ower On VM**, and so on . The actions displayed change based on the type of object you select. The button is dimmed when actions are not available for an object you select.
Constrained by	Sorts the objects in the chart based on a metric you select. For example, if you select CPU Demand, all the objects constrained by CPU demand are displayed in the chart. You can sort the chart based on options like: **CPU**, **CPU Demand**, **Memory**, **Memory Consumed**, and **vSphere Configuration Limit**.
Reset to initial object	Displays the original view of the chart.

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.

Option	Description
Show	Indicates whether the Capacity Utilization widget displays the capacity remaining or the workload balance for the objects in your environment. ■ Capacity Remaining. Displays a visual summary of the capacity resources used by your objects, and indicates why the objects are constrained. ■ Workload Balance. Displays a visual summary of the workload resources used by your objects, and indicates why the objects are constrained.
Select Object	Your inventory explorer where you can locate the object on which you are basing the data that appears in the widget.
Object Type	Select specific object types to see in the charts. Press Ctrl+click to select multiple object types. If you leave the object type deselected, you will see all base object children in the charts.

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.

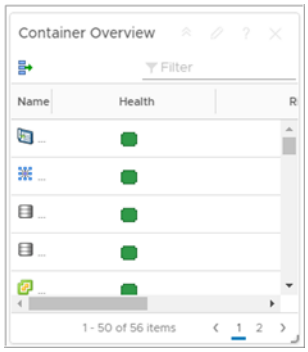

How the Container Overview Widget Works

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.

Container Overview Widget Toolbar Options

The toolbar at the top of the Container Overview widget 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.

The data grid provides information on which you can sort and filter.

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

To configure a widget, click the **Edit Widget** icon on the widget title bar.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
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 .
Object tree	The object tree appears when you select **Object** from the Mode option. You can filter the list of objects in the object data grid. You can select one or more object types and the data grid displays all objects from the types. For example, if you want to observe information about the VMs and vCenter Server in the inventory, click **Collapse All**, expand **Object Types** in the object tree, and select **Virtual Machine** and **vCenter Server**. As a result, the data grid shows only VMs and vCenter Server objects from the inventory. You can deselect adapter types when you click **Deselect All**.

Option	Description
Object data grid	**Note** The object data grid exists when you select **Object** from the **Mode** option. 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. When you click an object from the list it appears in the Selected Objects pane. You can select multiple objects from the data grid when you mark objects in the list and click the **Perform Multi-Select Interaction** toolbar icon. To deselect an object or objects, click the **Clear Selections** toolbar icon.
Selected Object	The Selected Object pane appears when you select **Object** from the Mode option. Object that is the basis for the widget data. You can add an object when you select it first from the object data grid. You can remove an object from the list when you select an object and click the **Delete Object** toolbar icon.
Selected Object Type	Selected Object Type appears when you select **Object Type** from the Mode option. Selecting this option shows the type of the objects to observe.
Object Type list	Selected Object Type exists when you select **Object Type** from the Mode option. By default, the list shows all available object types in the environment. You can select a type when you click a type in the list. 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. You can remove filtering when you click the plus sign in the drop-down menu.

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 Works

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. When you select an object in the Object List on the dashboard, the widget displays the policy associated with that object.

After you add the Current Policy widget to a dashboard, you click the pencil on the widget toolbar to edit the widget and configure the information to view in the widget. The changes that you make to the widget, including the Self Provider setting, and whether you select an object in the widget when you edit it, 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.

Where You Find the Current Policy Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The Current Policy widget includes toolbar options to collapse, edit, get help, and close the widget. To add the Current Policy widget to a dashboard, you create or edit a dashboard, click the widget in the widget list, and drag it to the dashboard workspace. After you add the widget to the dashboard, you configure the widget.

With the Current Policy widget configured, when you select an object on the dashboard, such as in the Object List widget, 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.

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The Current Policy widget requires you to either set the widget to be a Self Provider or to configure the widget interactions so that the widget receives the data required to indicate the policy that is applied to an object.

- To set the Current Policy widget as a self provider, you edit the widget configuration and select **Self Provider**.
- To have an object, such as the Object List widget, provide data to the Current Policy wizard on a dashboard, when you create or edit the dashboard, you click **Widget Interactions**, and select an object in the workspace to provide data to the Current Policy wizard.

See Widget Interactions.

Current Policy Widget Configuration and Data Grid Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.

Option	Description
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. For example, to view the policy applied to each object that you select in the Object List widget, for Self Provider you would select **Off**.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Per Page	Number of objects to view on each page.
Search	Locate data in the widget.
Policy	Operational policy applied to the object or object group.
Name	Object or object group name.
Description	Object or object group description.
Adapter Type	Adapter to which the object applies.
Object Type	Object type or object group type.
Policy	Name of the policy applied to the object or object group.
Creation Time	Date and time that the policy was created.
Maintenance Schedule	Date and time to perform maintenance tasks, if defined for the policy. vRealize Operations Manager does not collect metrics or calculate analytics during maintenance times.
Identifiers 1-5	Unique identifier for each object. These identifiers imply relationships between objects. ■ Identifier 1. Object name, which is the same as Name, and can include the full domain name. ■ Identifier 2. Object or object group identifier, including the type and number for each object such as a virtual machine, datacenter, host, and so on. ■ Identifier 3. Specific object identifier, or long identifier. ■ Identifier 4. Long identifier. ■ Identifier 5. IP address of the object.
Object Flag	Indicates the state of the object. For example: Normal.
Collection State	Indicates the state of vRealize Operations Manager collecting data from objects.
Collection Status	Indicates the status of the collection.

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 Works

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

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Data Collection Results Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.

Table 4-124. Data Collection Results Widget Configuration Options

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 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 .

Table 4-124. Data Collection Results Widget Configuration Options (continued)

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.
Config tab	Specifies self provider choice and selection of a resource instance.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the objects list.
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. If you select an object in the list, the object becomes the selected object for the widget. If you select an object in the list, the object becomes the selected object for the widget.
Per Page	Number of objects to view on each page.
Filter	Locate data in the widget.
Defaults tab	Specifies the default data collection action selected for each object type.
Object Types	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. If you select an object in the list, the object becomes the selected object for the widget. If you select an object in the list, the object becomes the selected object for the widget.
Default Data Collection Action	This panel is populated by the object you select in the object types list. You can select only one default data collection action for an object type.

Density Widget

The Density widget displays the density breakdown in charts for the past seven days for a specific resource.

The Density widget produces a graph depicting the concentration of objects in a particular state as a percentage. It compares the ideal consolidation ratio to the actual consolidation ratio. States that are displayed are Unknown state, Critical state, Immediate state, Warning state, and Normal state.

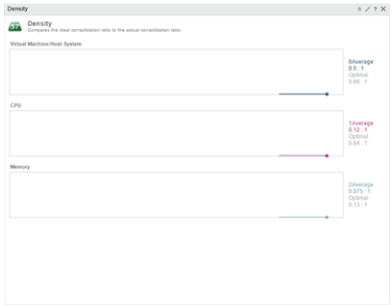

The Density widget configuration options are used to customize each instance of the widget that you add to your dashboards.

Where You Find the Density Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar.

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.
Selected Object	Object that is the basis for the widget data.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget.

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**.

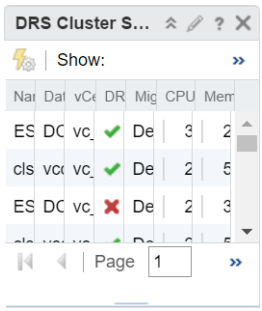

You edit a DRS Cluster Settings widget after you add it to a dashboard. To configure the widget, click the edit icon at the upper-right corner of the widget window. You can add the DRS Cluster Settings widget to one or more custom dashboards and configure it to display data that is important to different 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

DRS Cluster Settings Options and Configuration Options

The DRS Cluster Settings widget includes toolbar options, data grid options and configuration 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.

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.

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 .

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 Works

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

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Badge 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 set to Off. If the **Badge Mode** is on, only the badge appears.

Table 4-125. Efficiency Widget Configuration Options

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
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.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

Environment Widget

The Environment widget displays the resources for which vRealize Operations Manager 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.

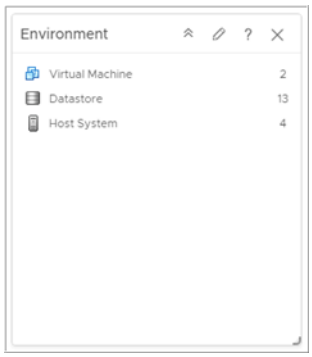

Where You Find the Environment Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The Weather Map widget provides for configurations options. To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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 .

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

Environment Overview Widget

The Environment Overview widget displays the health, risk, and efficiency of resources for a given object from the managed inventory.

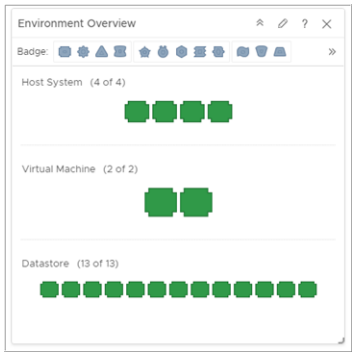

How the Environment Overview Widget Works

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.

Environment Overview Widget Toolbar Options

The toolbar at the top of the Environment Overview widget contains icons that you can use to get more information about badges.

Option	Description
Badge	You can select a badge for objects that appear in the widget. The tool tip of a badge shows the standard or custom name of the badge. You can add custom names to badges when you configure the widget by using the **Badge** tab.
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

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.

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.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .
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. Custom Label shows the custom name of a badge. You can use Custom Label to rename a badge. To rename a badge, double-click the badge and enter a name in a text box. To save the custom name, click **Update**. Custom Label is available only when custom metrics and badge customization are enabled.

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. If you select an object in the list, the object becomes the selected object for 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. The badge data grid shows custom badges and enables you to customize a badge for a custom metric. You can select a metric from the metrics tree and set the color of the badge. The Badge column contains the badge icons. The Metric column contains a custom metric that you can select from the metric tree. You can use Box Label text box to define a label of a badge. The tool tip description of a badge and tool tip of each object with this badge use the badge label. You can use Measurement Unit text box to define a measurement unit that is used in a tool tip description of each object. You can use Yellow Bound text box to define a value for which the badge is yellow. You can use Orange Bound text box to define a value for which the badge is orange.

Option	Description
	You can use Red Bound to define a value for which the badge is red. For example, if you want to observe the availability of a virtual machine and use the Health badge, you must select virtual machine as an object type, select the Health badge icon, search for availability in the metrics tree, and double-click it. You must define a meaningful label name and measurement unit to help you when observe the objects. You must specify different values for each color, for example -1 for yellow, 0 for orange, and 1 for red. You can use the metrics tree to select a metric that is specific for each object type. You can click **Select Object** to select specific metrics for an object. **Select Object** takes you to the object list data grid. The object list data grid displays all available objects in the environment and details about them.

Environment Status Widget

The Environment Status widget displays the statistics for the overall monitored environment.

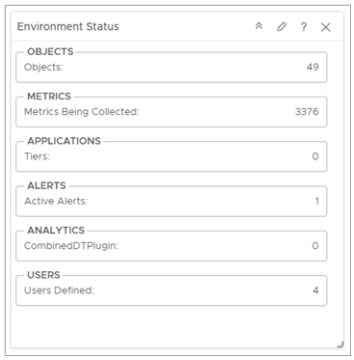

How the Environment Status Widget Works

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.

Environment Status Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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. 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.
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.

Option	Description
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**.
Select which tags to filter	You can select between different types of objects to observe.
	Use the **Collapse All** toolbar option to close all expanded tags and tag values.
	Use the **Deselect All** toolbar option to remove all filtering and view all objects in the widget.

Faults Widget

The Faults widget displays detailed information about faults experienced by an object

A fault score indicates the degree of problems that the object is experiencing. It includes events such as loss of redundancy in NICs or HBAs, memory checksum errors, HA failover problems and CIM events.

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

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Object List	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.

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 Works

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. On the menu, click **Dashboards** to see a list of dashboards in the left pane.

Where You Find the Forensics Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

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.

Option	Description
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .
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.
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.
Tag Tree	Filters the list of objects in the object list. You can select one or more object types and all objects from this type are displayed in the object list.
Object List	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. The objects show based on the selected tag. If no tag is selected, the list shows all objects in the system.
Metric Picker	Double-click the metrics to show in the widget.
Selected Object	Object that is the basis for the widget data.

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 **Geo** tab on the Inventory Explorer 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Geo Widget Toolbar Options

Option	Description
Zoom in	Zooms in on the map.
Zoom out	Zooms out on the map.

Geo Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Select which tags to filter	You can select between different types of objects to observe.
	Click the **Collapse All** toolbar option to close all expanded tags and tag values.
	Click the **Deselect All** toolbar option to remove all filtering and view all objects in the widget.

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.

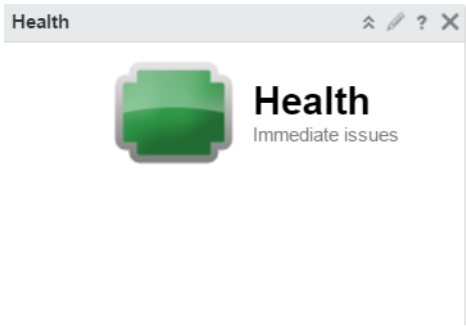

The state of the badge is based on your alert definitions. Click the badge to see the Summary tab for objects or groups configured in the widget. From the Summary tab you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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 is on, 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.
Badge Chart	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 set to Off. If the **Badge Mode** is on, only the badge appears.

Table 4-126. Heath Widget Configuration Options

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.

Table 4-126. Heath Widget Configuration Options (continued)

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.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

How the Health Chart Widget Works

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.

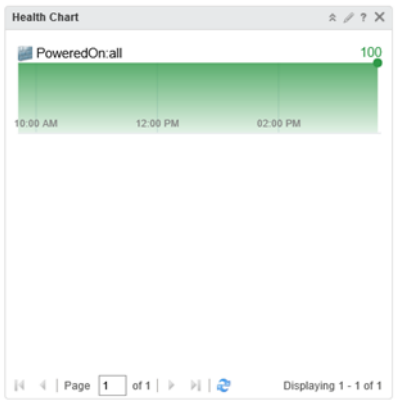

If the widget is configured to display Health, Risk, or Efficiency, the chart values are based on the generated alerts for the selected alert type for the selected objects.

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.

Where You Find the Health Chart Widget

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Health Chart Options

To view the value of the object at a particular time, hover your mouse over the chart. A date range and metric value tool tip appear.

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Mode	Determines if the widget displays data for the selected objects, child objects, or parent objects. If you select Children or Parents, the selected objects do not appear in the widget. Only the related objects.
Order By	Determines how the object charts appear in the widget. You can order them based on score or name, and in ascending or descending order.
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.
Period Length	Amount of time that is displayed in the chart.

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. If you apply custom colors, type the value in each box that is the highest or lowest value that should be that color. For example, if you select Custom, define the metric as Badge | Anomaly, and set Yellow Bound as 1, Orange as 10, and Red as 20, the charts display the changes from yellow to orange or red based on the anomaly metric values at each point in time.
Object Tag Tree	Object or object types for which to display charts. If you select a tag with more than one object, the widget displays charts for each object. If you select more than one tag, the widget displays charts only for the objects that are members of all the tags. If you select two tags an your widget does not display any charts, there were no common objects between the two tags.

Heat Map Widget

The Heat Map 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 Heat Map Widget Works

You can add the Heat Map widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The Heat Map 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 group the rectangles according to tag type and select the color range to use. 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.

When you point to a rectangle for an object, the widget shows the resource name, group-by values, and the current values of the two tracked attributes.

If you configure the Heat Map 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.

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.

You edit a Heat Map 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 Heat Map Widget

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The data that appears in the widget depends on how you configured it.

To configure the widget, click the **Edit Widget** icon on the widget title bar.

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.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget .
Configurations	List of saved Heat Map configuration options. You can create new 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 heat map.
Then by	Second-level grouping of the objects in the heat map.
Relational Grouping	After you select the Group by and Then by objects, select the **Relational Grouping** checkbox 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.
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 of the attributes that are defined for the resource kind.
Color by	An attribute to set the color of the rectangle for each resource.

Option	Description
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.
Filter	The widget shows information only for the that meet the filter conditions.

Mashup Chart Widget

The Mashup Chart widget shows disparate pieces of information for a resource. It shows a health chart, an anomaly count graph, and metric graphs for key performance indicators (KPIs).

How the Mashup Chart Widget Works

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.

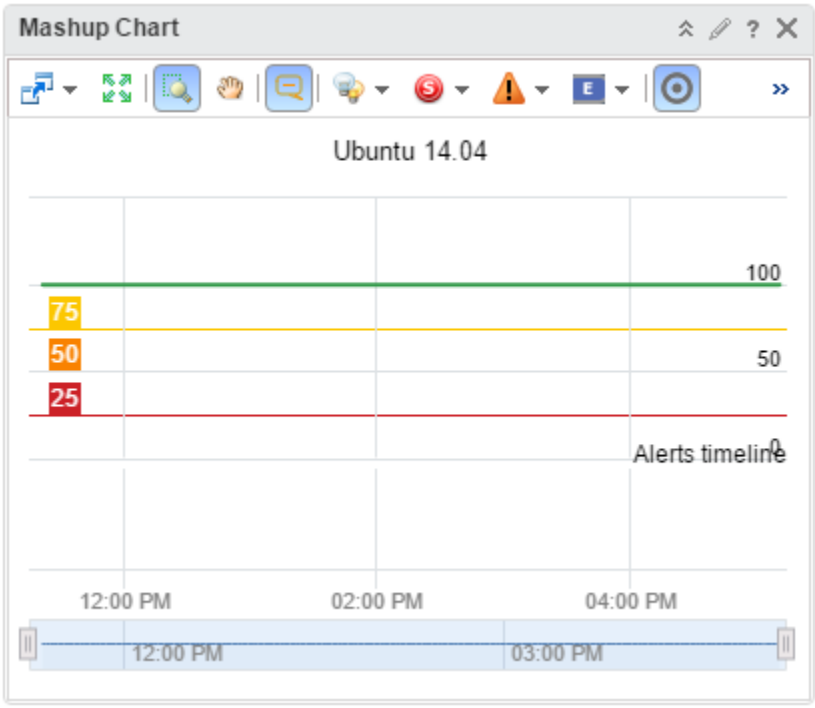

The Mashup Chart widget contains the following charts.

- 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.

- An Anomaly Count Graph for the object, which is similar to the anomaly graph that the cross-silo analysis feature generates. The graph shows the number of anomalies for the object and its children at the indicated time. For an application, it also shows the count for each tier in a stacked chart. A red line marks the noise threshold for the object. An anomaly count higher than this threshold indicates a 90 percent probability of a problem and triggers an early warning alert.
- Metric graphs for any or all of 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. Click the top left of the shaded area to see details about the anomaly.

The Anomaly Count Graph chart and 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. On the menu, click **Dashboards** to see a list of dashboards in the left pane.

The toolbar at the top of the Mashup Chart widget contains icons that you can use to change the view.

Mashup Chart Widget Configuration Options

You can customize the data that appears in the widget while you create the dashboard, or when the dashboard is displayed, by clicking the **Edit** icon in the title bar of 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.
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
Tag Tree	Filters the list of objects in the object list. You can select one or more object types and all objects from this type are displayed in the object list.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget. The objects show based on the selected tag. If no tag is selected, the list shows all objects in the system.

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.

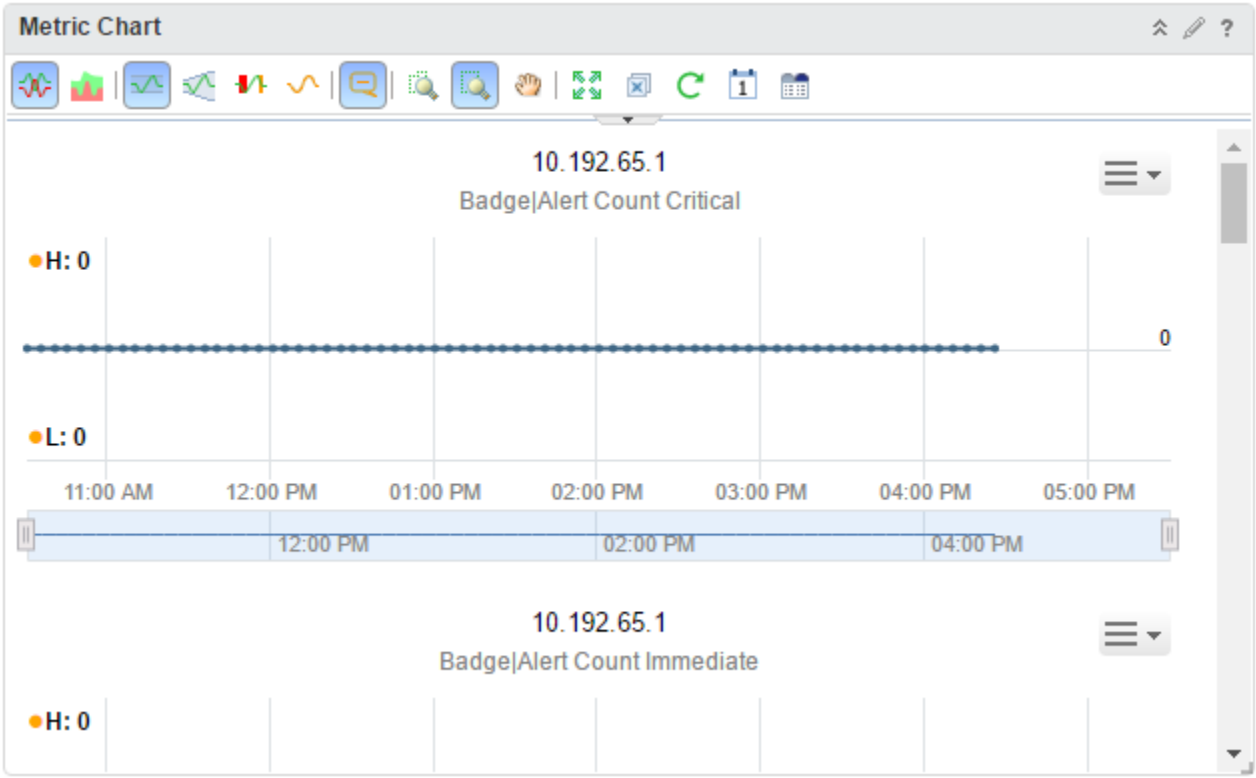

How the Metric Chart Widget Works

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. On the menu, click **Dashboards** to see a list of dashboards in the left pane.

The Metric Chart widget is also displayed on the Workload Utilization dashboard with the name Workload Trend.

Metric Chart Widget Toolbar Options

The toolbar at the top of the Metric Chart widget 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.
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.
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.

Option	Description
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.
Generate Dashboard	Saves the current charts as a dashboard.

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.
Select the units for the widget display	You can display the data with dots or as a percentage.
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.

Option	Description
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.

Metric Chart Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

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.
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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.
Object Tag Tree	Filters the list of objects in the object list. You can select one or more object types and all objects from this type are displayed in the object list.
Object List	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. The objects show based on the selected tag. If no tag is selected, the list shows all objects in the system.
Metric List	List of metrics available for the object selected in the object list. Double-click the metrics to show in the widget.
Selected Metric List	Objects and metrics that are displayed in the widget. The objects appear on the widget in the order they are presented in the list. To change the order of the displayed charts, reorder the list.

Metric Picker Widget

The Metric Picker widget displays a list of available metrics for a selected object.

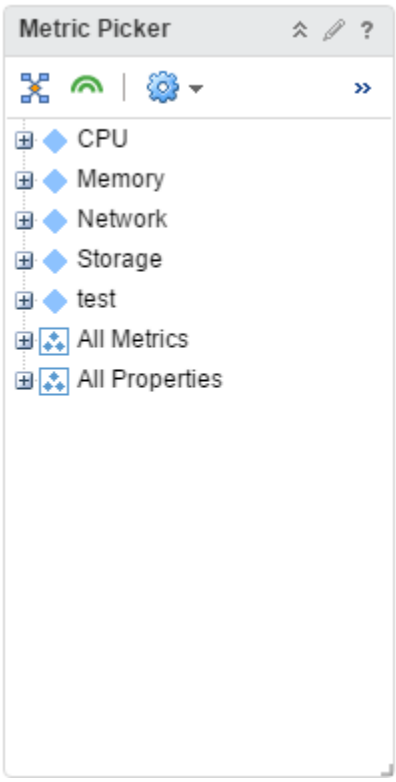

How the Metric Picker Widget Works

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 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 on any of your custom dashboards. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Metric Picker Widget Toolbar

The toolbar at the top of the Metric Picker widget 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.

Metric Picker Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

Table 4-127. Metric Picker Widget Configuration Options

Option	Action
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 .

Object List Widget

The Object List widget displays a list of the objects available in the environment.

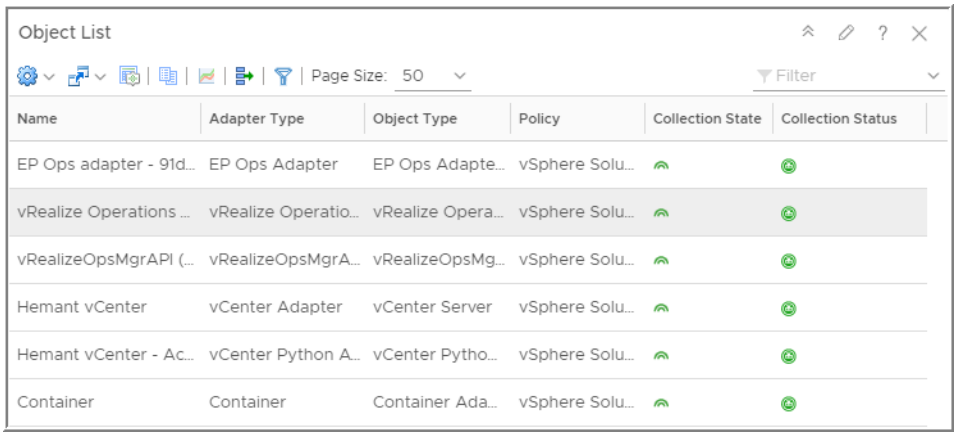

How the Object List Widget and Configuration Options Work

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.

Where You Find the Object List Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Object List Widget Toolbar and Datagrid Options

The Object List widget includes 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.
Filter	Locate data in the widget.

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
Adapter Type	Shows the adapter type for each object .

Option	Description
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 **Configuration > Inventory Explorer**.
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 **Configuration > Inventory Explorer**.
Internal ID	Unique number that vRealize Operations Manager uses to identiy the object internally. For example, the internal ID appears in log files used for troubleshooting.

Object List Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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 .
Mode	You can select the **Self**, **Children** , or **Parent** mode of the widget in the dashboard. For example, you can add two Object List widgets to a dashboard, with names Object List 1 and Object List 2. You can configure Object List 1 as a sender and Object List 2 as a receiver in the **Widget Interactions** option when you edit the dashboard. If Object List 2 is in self mode and you select an object from Object List 1, Object List 2 displays information only for the object that you selected. If you select parent mode for Object List 1 and children mode for Object List 2, the Object List 2 widget displays only children objects of an object that you select from Object List 1. For example, if you select host system from Object List 1, the Object List 2 widget displays all VMs on this host.
Auto Select First Row	Determines whether to start with the first row of data.

Option	Description
Select which tags to filter	Selects an object or objects from an object tree to observe. For example, to observe information about the VMs and vCenter Server in the inventory, you must click **Collapse All** and select **Virtual Machine** and **vCenter Server** under **Object Types** .
Additional Column	Adds columns with metrics that are specific for each object to the data grid. To add a metric, click **Pick Metrics** to go to the Pick Metrics with Object Type dialog box . You can explore available metrics for an object type and select a metric. ■ Object Types pane - Use to select the object type from the tree with object types. Metrics in Metric Picker Tree depend on your selection of an object type. ■ Adapter Type drop-down menu - Use to filter objects in the list based on an adapter that they use. All available adapter types are selected by default. You can select a concrete type using drop-down menu. You can select all adapter types using the close sign next to the drop-down. ■ Metric Picker - Use to select one or more metrics to observe. The metrics list is different for each object depending on its type and its instance. Each metric that you select is added to the Selected Metrics data grid. ■ Perform Multi-Select Interaction - Use to select several metrics from the metrics tree. ■ Select Object - Use to select an object to pick-up metrics. ■ Selected Metrics - Use to remove selected metric, sort the metrics, reorder them and manipulate the data grid columns.

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.

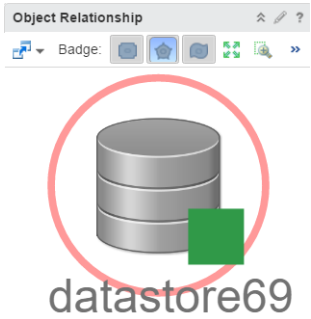

You edit an Object Relationship widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

Where You Find the Object Relationship Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Object Relationship Widget and Configuration Options

The Object Relationship widget includes 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.

Option	Description
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.
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.

The Object Relationship widget provides the following configuration options.

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
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.

Option	Description
Object Selection	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.
Select which tags to filter	Filter the parent and child objects that appear in the widget. If you select a tag, only parent and child objects that match your selection here appear in the widget. To show all of the parent and child objects of the selected object, do not select a tag value.

Object Relationship (Advanced) Widget

The Object Relationship (Advanced) widget displays the hierarchy tree for the selected object. It provides advanced configuration options. 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 (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.

Where You Find the Object Relationship (Advanced) Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Object Relationship (Advanced) Widget Toolbar and Configuration Options

The Object Relationship (Advanced) widget includes 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 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 select only 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.

Options	Description
Show values on point	Shows or hides the data point tooltips when you hover the mouse over a data point in the chart.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.
Zoom the view	Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section.
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.
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.
Pagination	Lets you select the number of parent or child objects to be displayed. The default value is 1-100.

The Object Relationship (Advanced) widget includes these configuration options.

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.

Option	Description
Object selection	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.
Select which tags to filter	Filter the parent and child objects that appear in the widget. If you select a tag, only parent and child objects that match your selection here appear in the widget. To show all of the parent and child objects of the selected object, do not select a tag value.

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) or 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 list of the widget configuration window.

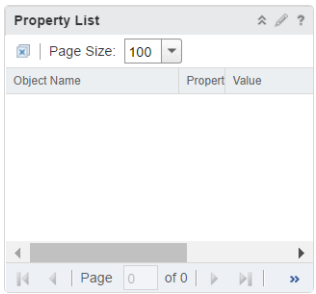

You edit a Property List widget after you add it to a dashboard. You can configure a widget to receive data from another widget by selecting **Off** for Self Provider mode. When the widget is not in Self Provider mode it displays a set of predefined properties and their values of an object that you select 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Property List Widget and Configuration Options

The Property List widget includes data grid options.

Option	Description
Object Name	Name of the object, whose properties you observe. You can sort the properties by object name. Click on an object name to open the Object Details page.
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.

The Property List widget includes configuration options.

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.
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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.

Option	Description
Objects	■ Objects tree Use to filter objects in the object list data grid. For example, you can expand **Object Types** and select **Virtual Machine** to observe only VMs from your inventory in the object list data grid. ■ Object list 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. If you select an object in the list, the object becomes the selected object for the widget.
Properties Tree	Double-click a property of the object selected from the Object list to observe in the widget.
Selected Object	Object that is the basis for the widget data.

Reclaimable Capacity Widget

The Reclaimable Capacity widget displays a percentage chart representing the amount of reclaimable waste for a specific resource which has consumers.

Where You Find the Reclaimable Capacity Widget

In the menu, click **Dashboard > Actions > Create Dashboard** to add a dashboard or click **Dashboard > Actions > Edit Dashboard** to edit the selected dashboard. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The data that appears in the widget depends on how you configured it.

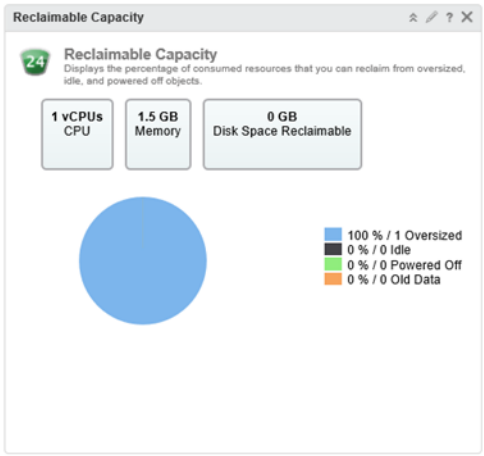

To configure the widget, click the **Edit Widget** icon on the widget title bar.

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.
Selected Object	Object that is the basis for the widget data.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

Where You Find the Recommended Actions Widget and Configuration Options

To access the Recommended Actions dashboard, in the menu, click **Home**, and then in left pane click **Recommended Actions**.

Recommended Actions Widget Options

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.

Table 4-128. 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.

The Recommended Actions widget includes a toolbar and a data grid that displays the alerts that triggered.

Table 4-129. Toolbar and Data Grid

Option	Description
Toolbar	The toolbar allows you to address an alert, and to filter the alert list. ■ 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. ■ Quick Filter (Alert)
Data Grid	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. ■ 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. ■ Alert Type.

Table 4-129. Toolbar and Data Grid (continued)

Option	Description
	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. For more information, see All Alerts.

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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 .

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 dashboard 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.

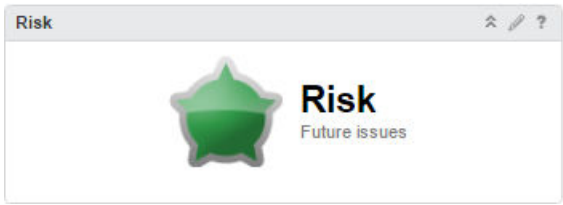

Where You Find the Risk Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

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.
Badge Chart	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 set to Off. If the **Badge Mode** is on, only the badge appears.

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 .

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.
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
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.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

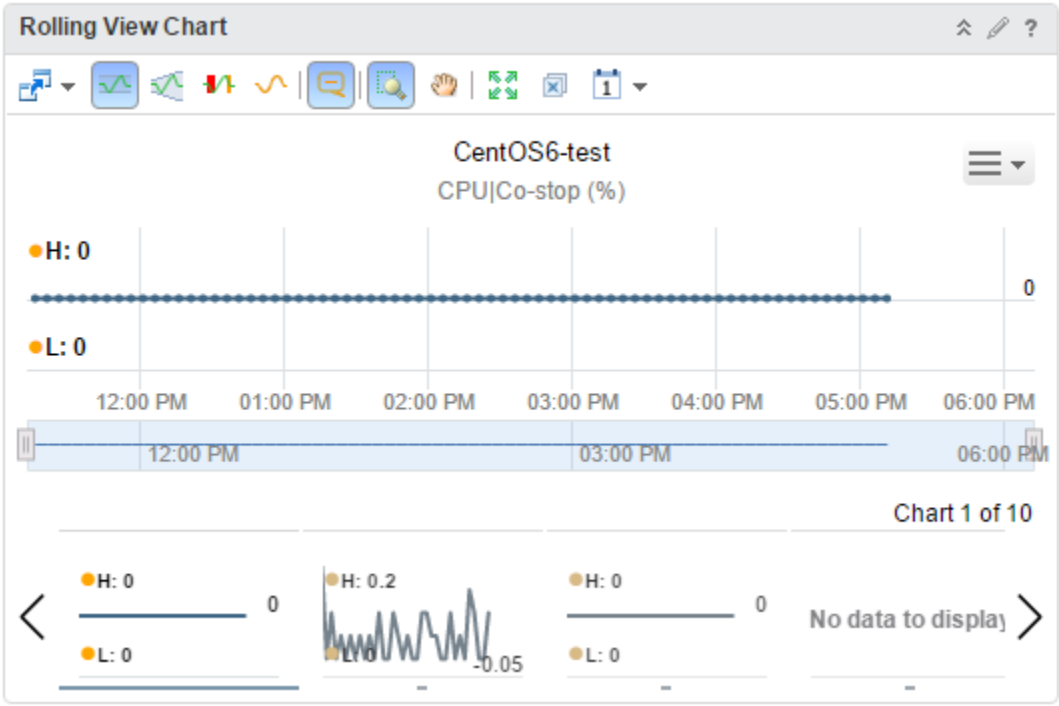

How the Rolling View Chart Widget Works

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. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Rolling View Chart Widget Toolbar

The toolbar at the top of the Rolling View Chart widget contains icons that you can use to change the view of the graphs.

Icon	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.

Icon	Description
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.

Rolling View Chart Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

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.

Option	Description
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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.
Auto Transition Interval	Time interval for a switch between charts in the widget.
Show Chart Toolbar	Determines whether the Toolbar options appear in the widget.
Tag Tree	Filters the list of objects in the object list. You can select one or more object types and all objects from this type are displayed in the object list.
Object List	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. The objects show based on the selected tag. If no tag is selected, the list shows all objects in the system.
Metric Picker	Double-click the metrics to show in the widget.
Selected Object	Object that is the basis for the widget data.

Scoreboard Widget

The Scoreboard widget shows the current value for each metric of objects that you select.

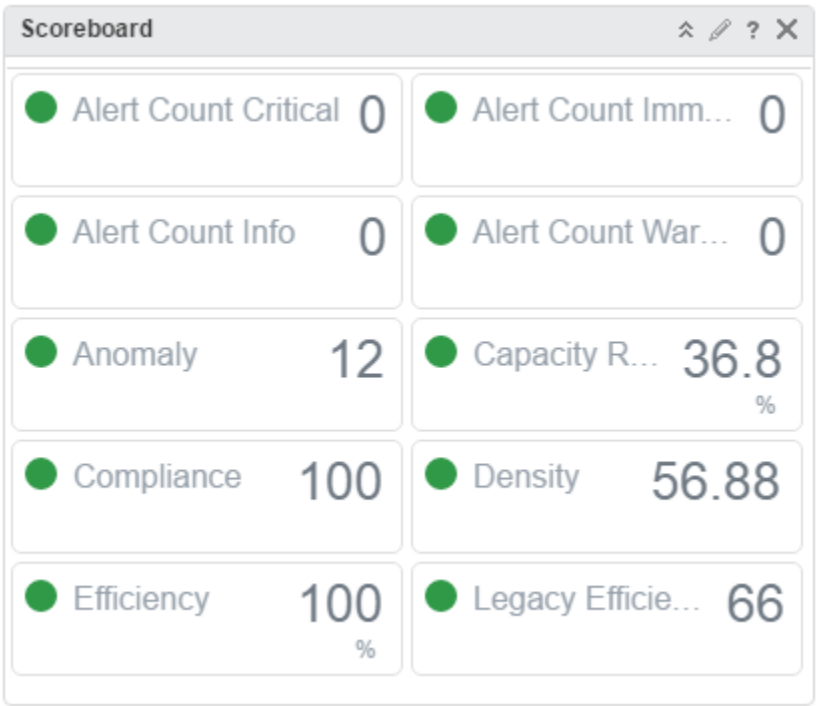

How the Scoreboard Widget Works

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.

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. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Scoreboard Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

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.

Option	Description
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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 Mode	Select a Fixed Size or Fixed View layout.
Box Height Box Columns	Use these menus to customize the size of the box for each object.
Visual Theme	Select a predefined visual style for each instance of the widget. The options are: **Original**, **Theme 1**, **Theme 2**, **Theme 3**, and **Theme 4**. The default style is Theme 2.
Label Size Value Size	Use these menus to customize the format of the scores that the widget displays.
Show Object Name	Select whether to display the object name.
Show Metric Name	Select whether to display the name of the metric in the widget. ■ **On**. The name of the metric you select is displayed in the widget. ■ **Off**. The name of the metric you select is not displayed in the widget.
Show Metric Unit	Select whether to display the metric unit in the widget. ■ **On**. The name of the metric you select is displayed in the widget. ■ **Off**. The name of the metric you select is not displayed in the widget.
Show Sparkline	Select whether to display the Sparkline chart for each metric. If you select the widget to display sparkline, you can select the time frame from the**Period Length** option that the chart includes.
Period Length	Select a length of time for statistic information that the sparkline chart shows.

Option	Description
Objects	Object that is the basis for the widget data. ■ Object tree You can filter the list of objects in the object data grid. You can select one or more object types and the data grid displays all objects from these types. For example, to observe information about the VMs and vCenter Server in the inventory, click **Collapse All**, expand **Object Types** in the object tree, and select **Virtual Machine** and **vCenter Server**. The data grid shows only VMs and vCenter Server objects from the inventory. To deselect adapter types, click **Deselect All**. ■ Object data grid Lists objects in your environment that you can search or sort by column so that you can locate the object to pick its metrics. When you click an object from the list, its metrics appear in the metric tree. You can select multiple objects from the data grid when you mark objects in the list and click the **Perform Multi-Select Interaction** toolbar icon. To deselect an object or objects, click the **Clear Selections** toolbar icon.
Object Types	List of available object types. Use to select an object type that is the basis for the metrics tree. You can select an object from an object type and to pick its metrics when you click **Select Object** toolbar icon from the metrics pane. The Select Object takes you to the list of objects from the selected object type. For example, you can select **Datacenter** from the Object Types data grid and click **Select Object** to display the list of datacenters in your environment.

Option	Description
Metric Tree	Shows available metrics of an object or object type that you select from the data grid. Use the metric tree to select a metric that is basis for the widget. The metric tree can show common metrics for several objects when you click the **Show common metrics** toolbar icon. To pick several metrics, select the metrics from the tree and click **Perform Multi-Select Interaction**. The **Select Object** toolbar icon appears when you use the **Object Types** tab.
Objects List	List of the objects and their metrics that the widget displays. Your selection of an object and a metric from the object data grid and metrics tree is propagated to the Object and Metric columns. You can use the **Box Label** text box to customize the label of each metric box on the widget. You can use the **Measurement Unit** text box to define a measurement unit of each metric. You can use the **Color Method** option to define a coloring criteria. To define values for the metric box colors, enter values in the text boxes. If you do not want to use color, select **None**. You can use the **Apply to All** toolbar icon to customize a metric box and apply the same customization to all metrics. For example, to select to observe a 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. To apply the same label and color criteria to all other selected metrics, select the metric and click **Apply To All**.

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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Scoreboard Health Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

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.
Image Type	Select an image type for the metrics.
Metric	Select the default or custom metric.

Option	Description
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.
Object Tree	Use to filter the objects in the object list. For example, you can expand **Object Types** and select **Virtual Machine** to observe only the VMs in your environment.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget. Use Perform-MultiSelect Interaction to select more than one object at a time from the data grid . You must mark the objects and click **Perform-MultiSelect Interaction**.
Selected Objects	Object that is the basis for the widget data. Your selection of an object from the Object List option is propagated to the list of Selected Objects.

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.

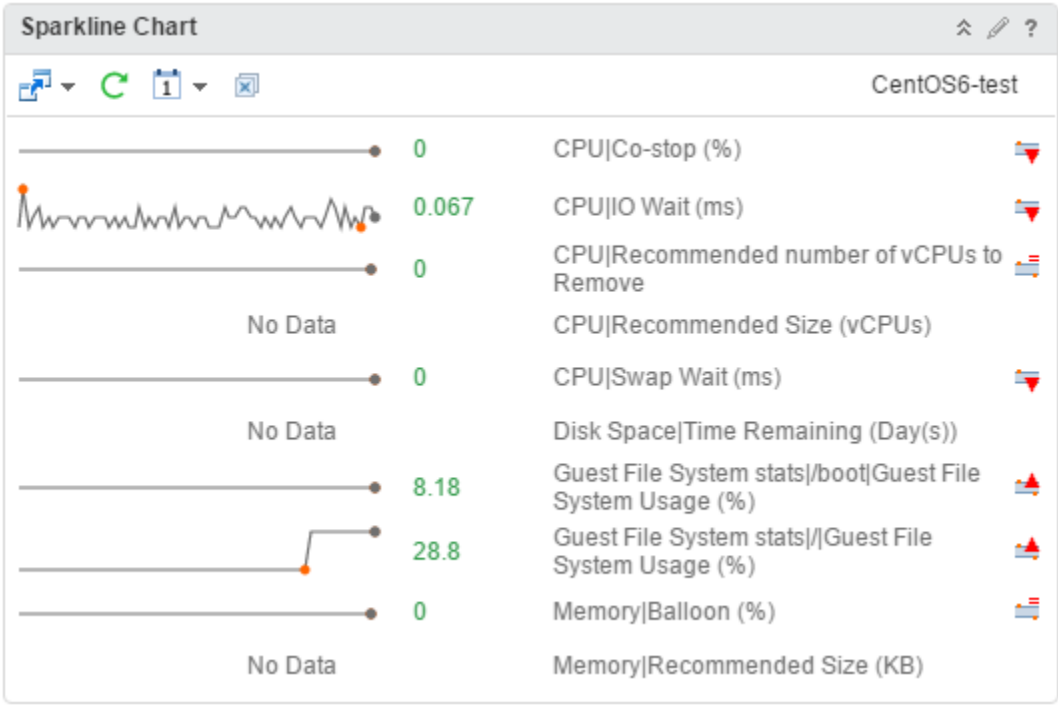

How the Sparkline Chart Widget Works

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

The toolbar at the top of the Sparkline Chart widget contains icons that you can use to change the view of the graphs.

Icon	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.
Remove All	Removes all graphs.

Sparkline Chart Widget Configuration Options

To configure a widget, click the **Edit** icon on the widget title bar.

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.
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.
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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.

Option	Description
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.
Objects	You can select metrics for specific objects during widget configuration. You can select one or more tag values to filter the objects to appear in the pane that lists the objects. You can use icons on the toolbar at the top of the list to collapse and deselect all the tags in the list.
	In the pane that lists the objects, use the toolbar options to select one or more objects. ■ To clear all your selections, click the **Clear Selection** icon. ■ To select multiple objects, click the**Perform Multi-Select Interaction** icon. ■ To set the number of objects to display in the pane, select a value in the **Page Size** field. ■ To search for an object, enter all or part of the object name in the **Filter** text box. The corresponding metrics for the selected object appear in the pane that lists the metrics.
	In the pane that lists the metrics, use the toolbar options to select the metrics to show in the widget. ■ To select multiple metrics, click the**Perform Multi-Select Interaction** icon on the toolbar at the top of the pane. ■ To list the metrics that are common to multiple selected objects, click the **Show common metrics** icon on the toolbar. ■ To view an object , click the **Select Object** icon on the toolbar. ■ To search for a specific metric, enter all or part of the metric name in the **Filter** text box.

Option	Description
	You can configure the metrics for the selected objects. Set values for each metric in the pane that displays the selected metrics. To enter a value, point to the text box under the column heading, double-click within the text box, and enter the value. ■ **Box Label**. A label for the metric. ■ **Measurement Unit**. The measurement unit that appears after the metric value. ■ **Color Method**.To define values for the metric box colors, enter values in the text boxes. Select **Custom** to set the color boundaries. If you do not want to use color, select **None**. You can manage the metrics in the pane that displays the metric. ■ To select all the metrics in the list, click the **Select All**icon on the toolbar at the top of the pane. ■ To remove all the metrics from the list, click the **Clear Selections** icon on the toolbar at the top of the pane. ■ To apply settings from one metric to all the metrics in the list, select the metric and click the **Apply To All** icon on the toolbar at the top of pane.
Object Types	You can select metrics for specific object types during widget configuration. This option is useful if specific objects are not currently available. To select an object type, use the icons on the toolbar. ■ To search for a specific adapter, you can enter the name of the adapter in the **Adapter Type** text box. ■ To search for an object, you can enter all or part of the object type name in the **Filter** text box. The metrics for the object type appear in the pane that lists the metrics. You can select multiple metrics.

Option	Description
	In the pane that lists the metrics, use the toolbar options to select the metrics to show in the widget. ■ To select multiple metrics, click the **Perform Multi-Select Interaction** icon on the toolbar. ■ To list the metrics that are common to multiple selected object types, click the **Show common metrics** icon on the toolbar. ■ To select a specific metric that is specific to an object, click the **Select Object** icon on the toolbar. ■ To search for a specific metric, enter all or part of the metric name in the **Filter** text box.
	You can configure the metrics for the selected object types. Set values for each metric in the pane that displays the selected metrics. To enter a value, point to the text box under the column heading, double-click in the text box, and enter the value. ■ **Box Label**. A label for the metric. ■ **Measurement Unit**. The measurement unit that appears after the metric value. ■ **Color Method**. To define values for the metric box colors, enter values in the text boxes. Select **Custom** to set the color boundaries. If you do not want to use a color, select **None**. You can manage the metrics in the pane that displays the metric. ■ To select all the metrics in the list, click the **Select All**icon on the toolbar. ■ To remove all the metrics from the list, click the **Clear Selections** icon on the toolbar. ■ To apply settings from one metric to all the metrics in the list, select the metric and click the **Apply To All** icon on the toolbar.

Stress Widget

The Stress widget displays a weather map of the average stress for a specific resource for a time interval.

Stress is defined as when demand on resources chronically consumes most of your capacity. The Stress zone helps identify hosts and virtual machines that do not have enough resources allocated.

Where You Find the Stress Widget and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

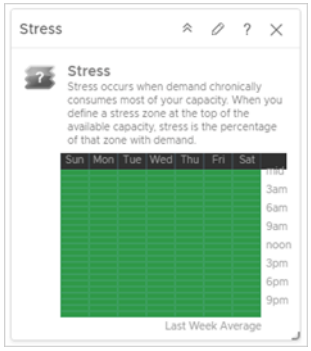

To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

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 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.

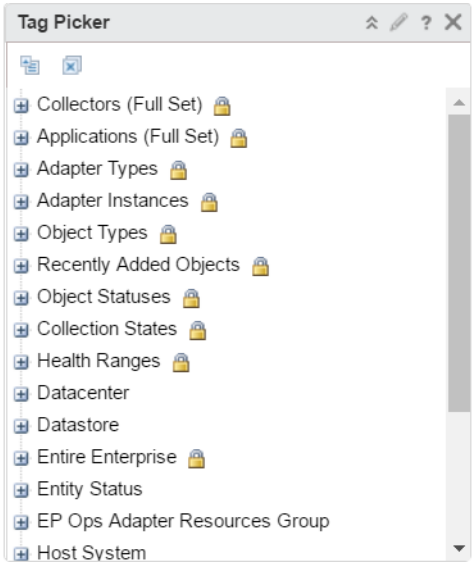

You edit a Tag Picker widget after you add it to a dashboard. To configure the widget, click the pencil in the upper-right of the widget window. You can configure the Tag Picker widget to send information to another widget on the same dashboard or on another dashboard. To set a receiver widget that is 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Tag Picker Widget and Configuration Options

The Tag Picker widget includes 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.

To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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 .

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 specify the path to the directory containing the text files. For example, if you configure the vRealize Business for Cloud adapter, relevant data is displayed in the Business Management dashboard.

The Text Display widget can display web sites that use the HTTPS protocol. The behavior of the Text Display widget with web sites that use HTTP, depends on the individual settings of the web sites.

How the Text Display Widget Configuration Options Work

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.

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 Text Display Widget Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Table 4-130. Text Display Widget Configuration Options

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 .
View mode	Display text in text or HTML format.
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 score by resource type based on historical data for the pattern of use for the resource type. You can use the time remaining score to plan provisioning of physical or virtual resources for the object or rebalance the workload in your virtual infrastructure.

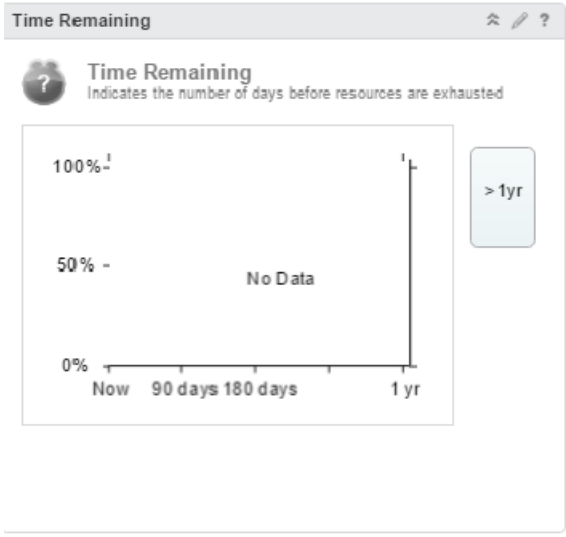

Where You Find the Time Remaining Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

Table 4-131.

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.

Table 4-131. (continued)

Option	Description
Selected Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Objects List	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. If you select an object in the list, the object becomes the selected object for the widget.

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.

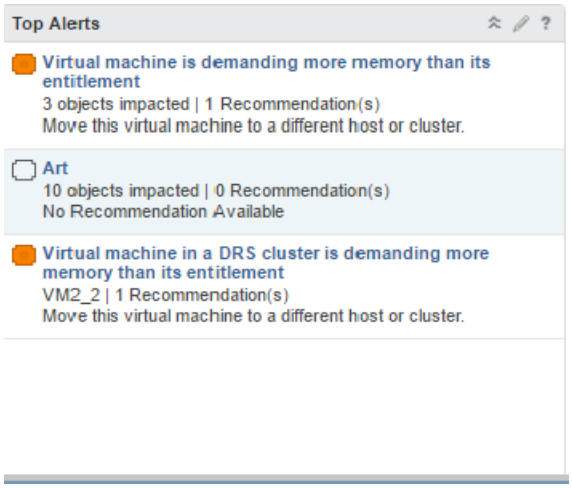

Where You Find the Top Alerts and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Top Alerts Data and Configuration Options

The top alerts include 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.

Table 4-132. Top Alerts Widget Options

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.

Table 4-133. Top Alerts Configuration Options

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.
Show Alerts On	Select one of the following options to specify the relationship of the objects that populate the widget data to the selected object. ■ Selected Object. The widget data is based only on the selected object. ■ Descendants Only. The widget data is based only on the descendant objects, not the selected object. ■ Both. The widget data includes both the selected object and the descendant objects.
Impacted 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.

Table 4-133. Top Alerts Configuration Options (continued)

Option	Description
Object	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Objects List	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. If you select an object in the list, the object becomes the selected object 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.

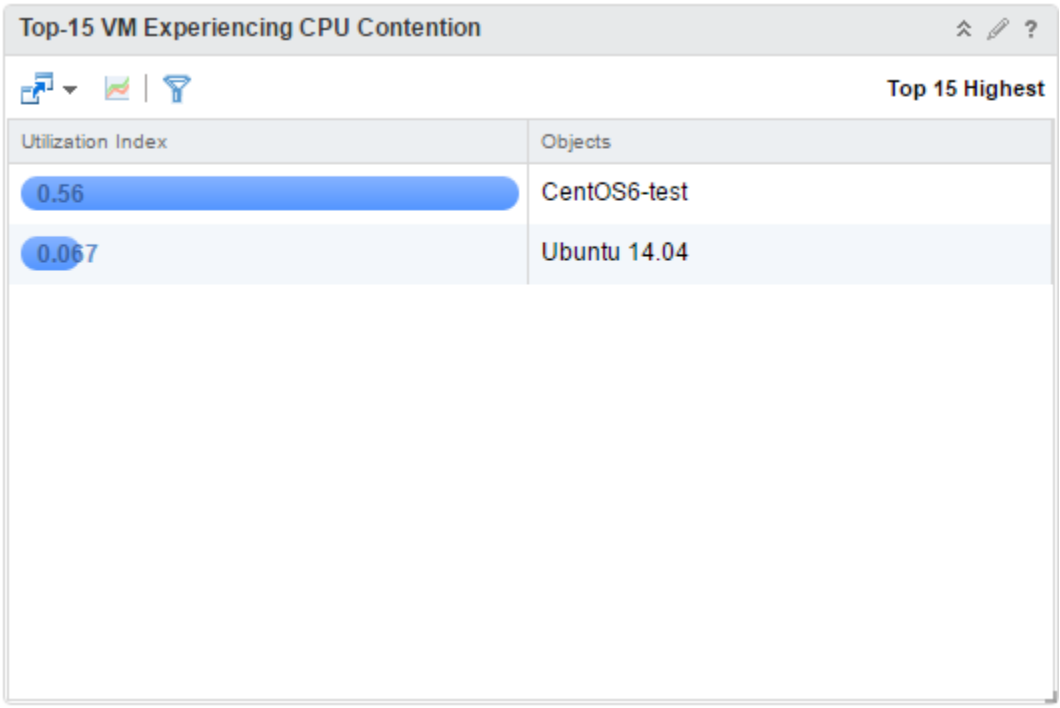

How the Top-N Widget Works

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. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Top-N Widget Toolbar

The toolbar at the top of the Top-N widget 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.
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

To configure a widget, click the **Edit** icon on the widget title bar.

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.
Image Redraw Rate	Set the redraw rate.

Option	Description
Period Length	Use the **Range** menu to select a range of time for which to display data. Use the **From** and **To** menus to select a specific start and stop date and time period. **Note** If you select **Current Value** as the range, the result is based on the last data collected. For any other range you select, the result is based on the aggregated value.
Application Health and Performance	Available when you use the **Tag** tab. ■ 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	Available when you use the **Tag** tab. Select the criteria for analysis of the alerts.
Metric Analysis	Available when you use the **Metric** tab ■ 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.
Bars Count	Select the number of top results.
Depth	Select the number of the child objects.
Filter old metrics	Select or deselect whether the analysis includes old metric values.
Selected Object	Object that is the basis for the widget data. The object that you select from the Objects data grid when you expand **Objects** is propagated to the text box.
Selected Object Type	The object type or types that you select from the Object Types data grid. Click the **Clear Selection** toolbar icon from the Object Type pane to clear the text box.

Option	Description
Tag	■ 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. If you select an object in the list, the object becomes the selected object for the widget. ■ Tag Picker. Use the Objects tag to select objects that are the basis for the widget. For example, you can click **Collapse All**, expand **Object Types** and select **Datacenter** and **Datastore** from the tag tree to observe data center and datastore objects from your inventory.
Metric	■ Tag tree Select the object tag that is the basis for the widget. For example, you can expand **Object Types** and select **Host System** to observe a metric analysis of the hosts in your environment. ■ Object type data grid Select one or more object types that are the basis for the widget. For example, you can select Virtual Machine and Compute Resource from the data grid and pick a common metric for both object types for analysis. The object types that you select from the data grid are propagated to the **Selected Object Type** text box. ■ Metric tree Select a metric that is the basis for the analysis that the widget shows. You can select a common metric or a metric that is specific for each object. To select a metric, first select an object type or object types from the data grid. For example, you can select **Virtual Machine** and **Datacenter** from the object types list and click **Show common metrics** to select a common metric for a VM and data center. You can click **Select Object** to select an object and pick a specific metric.

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.

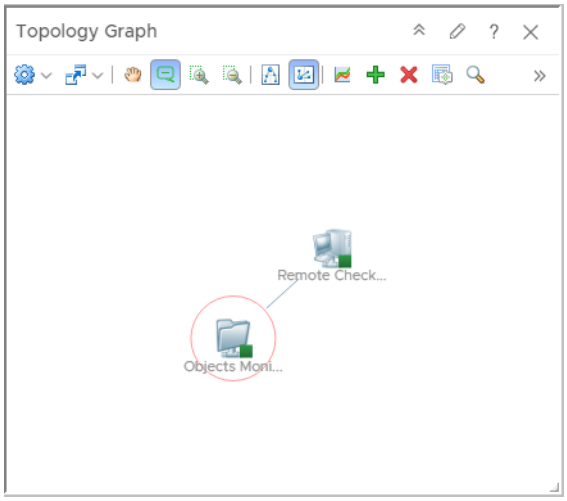

How the Topology Graph Widget and Configuration Options Work

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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

Topology Graph Widget 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

To configure a widget, click the **Edit Widget** icon on the widget title bar.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
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. Use **Show All Path** 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 Path 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 .
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 .
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.

Option	Description
Metric Configuration	Specifies a list with attributes to display, when the information is based on the interaction with another widget. To add a resource interaction XML file through the CLI directory, seeAdd 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.
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.
Selected Object	Object that is the basis for 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.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget.
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.
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.

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 properties, metrics, 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.

Where You Find the View Widget and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The View widget toolbar depends on the displayed view type. You can export the view as a CSV file for any view type.

View Widget Configuration Options

To configure a widget, click the **Edit Widget** icon on the widget title bar.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
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.
Select Object	Object that is the basis for the widget data.
Views	List of defined views, available for the selected resource. You can create, edit, delete, clone, export, and import views directly from the View widget configuration options.

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.

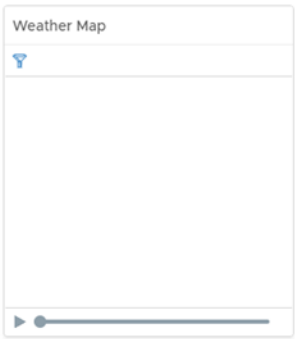

Watching how the map changes can help you understand how the performance of the metric varies over time for different resources. You can start or stop the display using the **Pause** and **Play** options at the bottom of the map. You can move the slider forwards or backwards to a specific frame in the 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 and Configuration Options

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. In the Dashboard workspace, on the left, click **Widget List**, and drag a widget to the right pane of the dashboard. On the title bar of the selected widget, click the **Edit Widget** icon to access the configuration options.

The toolbar at the top of the Weather Map widget contains the icons that you can use to view the graph.

Table 4-134. Metric Weather Map Widget Toolbar Icons

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.

The Weather Map widget provides for configurations options. To configure a widget, click the **Edit** icon on the widget title bar. For more information on creating and configuring dashboards, see Create and Configure Dashboards.

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 .
Image 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.

Option	Description
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. If you set a minimum or maximum value, any metric at or beyond that value appears in the end color.
Selected Object Type	Object that is the basis for the widget data. This text box is populated by the object you select in the Objects list.
Tag Tree	Filters the list of objects in the object list. You can select one or more object types and all objects from this type are displayed in the object list.
Object List	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. The objects show based on the selected tag. If no tag is selected, the list shows all objects in the system.
Metric Picker	Double-click the metrics to show in the widget.

Workload Widget

The Workload widget displays data a 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 and Configuration Options

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The data that appears in the widget depends on how you configured it. To configure the widget, click the **Edit Widget** icon on the title bar to configure the settings.

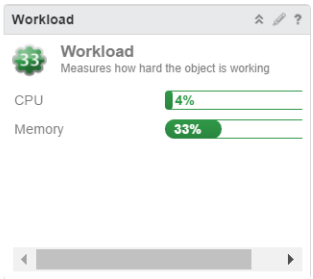

About Datastore Metrics for Virtual SAN

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.

The Workload Widget provides the following configuration options.

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.
Selected Object	Object that is the basis for the widget data.
Object List	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. If you select an object in the list, the object becomes the selected object for the widget.

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 4-135. Menu Options

Menu	Description
All Dashboards	Lists the dashboards that are enabled. You can use this menu for a quick navigation through your dashboards. When you navigate to a dashboard using the **All Dashboards** option, the dashboard is listed in the left pane of the Dashboards page.
Actions	Available dashboard actions, such as create, edit, delete, and set as default. These actions are applied directly to the dashboard that you are on.

Types Of Dashboards

You can use the predefined dashboards or create your own custom dashboard in vRealize Operations Manager.

Custom Dashboards

vRealize Operations Manager has predefined dashboards. You can also create dashboards that meet your environment needs.

To manage your dashboards, in the menu, click **Dashboards**.

Depending on your access rights, you can add, delete, and arrange widgets on your dashboards, clone and create dashboards, import or export dashboards from other instances, edit widget configuration options, and configure widget interactions.

Table 4-136. Dashboards Options

Option	Description	Usage
Save as Template	Contains all the information in a dashboard definition.	You can use any dashboard to create a template.
Export Dashboard	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.

Table 4-136. Dashboards Options (continued)

Option	Description	Usage
Import Dashboard	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.
Remove Dashboard(s) from Home	Removes a dashboard from the vRealize Operations Manager home page.	You can add any dashboard to the vRealize Operations Manager home page.
Reorder/Autoswitch Dashboards	Changes the order of the dashboard tabs on vRealize Operations Manager home page.	You can configure vRealize Operations Manager to switch from one dashboard to another.
Manage Summary Dashboards	Provides you with an overview of the state of the selected object, group, or application.	You can change the **Summary** tab with a dashboard to get information specific to your needs.
Manage Tab Groups	Groups dashboards in folders.	You can create dashboard folders to group the dashboards in a way that is meaningful to you.
Share Dashboards	Makes a dashboard available to other users or user groups.	You can share a dashboard or dashboard template with one or more user groups.

The dashboard list depends on your access rights.

Predefined Dashboards

vRealize Operations Manager has predefined dashboards that address several key questions including how you can troubleshoot your VMs, the workload distribution of your hosts, clusters, and datastores, the capacity of your data center, and information about the VMs. You can also log details.

In the menu, click **Dashboards**, and then in the left pane you can access a list of predefined dashboards. To enable dashboards to appear in the left pane of the Dashboards page, in the menu, click **Dashboards**, and then from the **All Dashboards** drop-down menu select the check-box against the required dashboard.

The default dashboard that appears when you click **Dashboards** in the menu is the **Getting Started** dashboard. You can close a dashboard from the left pane by selecting the dashboard and clicking the **X** icon. The dashboard you last opened is displayed the next time you navigate to **Dashboards** in the menu. If there is only one dashboard left in the left pane, you cannot close it.

You can access the following predefined dashboards from the left pane in the Home page.

- Recommended Actions
- Operations Overview
- Capacity Overview
- Workload Balance

The following predefined dashboards can be accessed by clicking **Dashboards** in the menu, and then clicking **All Dashboards**:

- Capacity and Utilization
  - Capacity Overview
  - Capacity Reclaimable
  - Cluster Utilization
  - Datastore Utilization
  - Heavy Hitter VMs
  - Host Utilization
  - VM Utilization
  - vSAN Capacity Overview
- Configuration and Compliance
  - Cluster Configuration
  - Host Configuration
  - Network Configuration
  - VM Configuration
  - vSphere Hardening Compliance
- Operations
  - Datastore Usage Overview
  - Host Usage Overview
  - Operations Overview
  - Optimize vSAN Deployments
  - vSAN Operations Overview
- Performance Troubleshooting
  - Troubleshoot a Cluster
  - Troubleshoot a Datastore
  - Troubleshoot a Host
  - Troubleshoot a VM
  - Troubleshoot vSAN
  - Troubleshoot with Logs
- vRealize Automation
  - vRealize Automation Environment Overview

    - vRealize Automation Top-N
- vRealize Operations
    - Self Cluster Statistics
    - Self Health
    - Self Performance Details
    - Self Services Communications
    - Self Services Summary
    - Self Troubleshooting
    - vCenter Adapter Details
- Getting Started

Getting Started Dashboard

The Getting Started dashboard is a guide to answering the most frequent questions of your IT staff. The dashboard breaks tasks into broad categories including Capacity and Utilization, Configuration and Compliance, Operations, Performance Troubleshooting, and Workload Balance.

Using each of these categories you can drill down to the specific use cases and problems you are trying to solve. Each problem statement is associated with a predefined dashboard that you can access through this page. To view a dashboard, click the dashboard name listed on the right side of the Getting Started dashboard.

Capacity and Utilization Dashboards

The dashboards in the Capacity and Utilization category cater to the teams responsible for tracking the utilization of the provisioned capacity in their virtual infrastructure. The dashboards within this category allow you to take capacity procurement decisions, reduce wastage through reclamation, and track usage trends to avoid performance problems due to capacity shortfalls.

Key questions these dashboards help you answer are as follows:

- How much capacity exists, how much is used, and the usage trends for a specific vCenter, data center, or cluster?
- How much disk, vCPU, or memory you can reclaim from large VMs in your environment to reduce wastage and improve performance?
- Which clusters have the highest resource demands?
- Which hosts are being heavily utilized and why?
- Which datastores are running out of disk space and who are the top consumers?
- The storage capacity and utilization of your vSAN environment along with the savings achieved by enabling deduplication and compression.

Capacity Overview Dashboard

The Capacity Overview Dashboard provides you with a summary of the total physical capacity available across all your environments being monitored by vRealize Operations Manager. The

dashboard provides a summary of CPU, memory, and storage capacity provisioned along with the resource reclamation opportunities available in those environments.

Since capacity decisions are mostly tied to logical resource groups, the Capacity Overview dashboard allow you to assess capacity and utilization at each resource group level such as vCenter, datacenter, custom datacenter, or vSphere cluster. You can quickly select an object and view the total capacity and used capacity of the object to understand the current capacity situation. Capacity planning requires that you have visibility into the historical trends and future forecasts. The trend views within the dashboard provide you with this information to predict how soon you will run out of capacity.

If you plan to report the current capacity situation to others within your organization, you can edit the Cluster Capacity Details widget on this dashboard and export this as a report for sharing purposes.

You can use the dashboard widgets in several ways.

- **Total Environment Capacity**: 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.
- **Total Reclamation Opportunity**: Use this widget to view the reclaimable resources in the environment.
- **Total Capacity**: Use this widget to view the total physical capacity of the environment that includes capacity assigned as High Availability (HA). The actual capacity is less than the total capacity displayed when you consider HA and a buffer.
- **Used Capacity**: Use this widget to view the capacity that has been used in your environment.
- **Memory Capacity Utilization Trend (TB)**: Use this widget to view the overall memory capacity trend. This widget displays the total physical resources you have. The physical resources include a HA buffer and a utilization buffer. This widget also displays the total memory you have allocated to VMs. If the number is close to the total physical capacity, the VMs may contend for memory. Ensure that the contention level is lower than what you promise to your customers. The chart also includes the actual utilization of memory capacity. The actual utilization is based on the active memory and hence it tends to be lower, as VMs do not normally access most of their RAM at any given moment.
- **CPU Capacity Utilization Trend (GHz)**: Use this widget to view the overall CPU capacity trend. This widget displays the total physical resources you have. The physical resources include a HA buffer and a utilization buffer, which reflects the total capacity. This widget also displays the total CPU capacity you have allocated to VMs. If the number is close to the total physical capacity, the VMs may contend for CPU. Ensure that the contention level is lower than what you promise to your customers. The chart also includes the actual utilization of

CPU. The actual utilization is based on the CPU demand counter, which takes into account the CPU used to perform I/O on behalf of the VM. The ESXi host performs storage I/O and network I/O on behalf of the VM, and this may be performed on a core that is different from the one on which the VM runs. As a result, CPU demand is a more accurate reflection of the VM CPU usage.

- **Disk Space Capacity Utilization Trend**: Use this widget to view the amount of disk space allocated to a VM and the amount that is actually used. This information is helpful when you plan for thin provisioning.
- **Cluster Capacity Details**: Use this widget to view the capacity of each cluster in the environment. You can view details such as the number of VMs, hosts, datastores, and CPUs in each cluster. You can also view details such as the total CPU capacity and the provisioned CPU capacity, the total memory, and the provisioned memory in each cluster.

Capacity Reclaimable Dashboard
The Capacity Reclaimable dashboard provides you with a quick view of resource reclamation opportunities within your virtual infrastructure.

This dashboard is focused on improving the efficiency of your environment by reducing the wastage of resources. This wastage is usually caused by idle or powered off virtual machines. Another biggest contributor to this wastage is oversized virtual machines.

This dashboard allows you to select an environment and quickly view the amount of capacity that can be reclaimed from the environment in the form of reclaimable CPU, memory, and disk space.

The dashboard lists all the virtual machines running on old snapshots and all the virtual machines that are powered off. You can reclaim storage by deleting the old snapshots on them or by deleting unwanted virtual machines. You can take these actions from the view by using the actions framework available withinvRealize Operations Manager.

The dashboard provides you with best practices about reclaiming CPU and Memory from large virtual machines in your environment. Since large and oversized virtual machines can increase contention between VMs, you can use the phased approach of using aggressive or conservative reclamation techniques to fix the size of your virtual machine.
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 Configuring and Using Workload Balance. If all the clusters in a given environment display the same pattern, you might have to add new capacity to cater to the increase in demand.

Datastore Utilization Dashboard

The Datastore Utilization dashboard helps you identify storage provisioning and utilization patterns in a virtual infrastructure.

As a best practice, ensure that the datastores are of standard size, to manage storage in your virtual environments. The heat map on this dashboard displays all the datastores monitored by vRealize Operations Manager and groups them by clusters.

The dashboard uses colors to depict the utilization pattern of the datastores. Grey represents an underutilized datastore, red represents a datastore that has run out of disk space, and green represents an optimally used datastore. You can select a datastore from the dashboard to see the past utilization trends and forecasted usage. The dashboard lists all the VMs that run on the selected datastore. You can reclaim storage used by large VM snapshots or powered off VMs.

You can use the vRealize Operations Manager action framework to reclaim resources by deleting the snapshots or unwanted powered off VMs.

- **Datastore Capacity and Utilization**: Use this widget to find out which datastores are overused and which ones are underused. You can also find out whether the datastores are of equal size. When you select a datastore from this widget, the dashboard is automatically populated with the relevant data.
- **VMs in the Selected Datastore**: Use this widget to view a list of VMs based on the datastore you select. You can also view relevant details such as whether the VMs are powered on and the size of the snapshot if any.
- **Usage Trend of Selected Datastore**: Use this widget to find out the trends in capacity used by a selected datastore as against the total capacity available.
- **All Shared Datastores in the Environment**: Use this widget to view a list of datastores that are shared in your environment. The information displayed in this widget helps you make an informed decision about whether you have to rebalance the capacity of the datastores based on usage.

Heavy Hitter VMs

The Heavy Hitter VMs dashboard helps you identify virtual machines which are consistently consuming a large amount of resources from your virtual infrastructure. In heavily over-provisioned environments, this might create resource bottlenecks resulting in potential performance issues.

You can use this dashboard to identify the resource utilization trends of each of your vSphere clusters. With the utilization trends, you can also view a list of VMs within those clusters based on their resource demands from the CPU, memory, disk, and network within your environment. You can also analyze the workload pattern of these VMs over the past week to identify heavy hitter VMs which might be running a sustained, heavy workload that is measured over a day, or bursty workloads that is measured using peak demand.

You can export a list of offenders and take appropriate action to distribute this demand and reduce potential bottlenecks.

You can use the dashboard widgets in several ways.

- **Select a Cluster**: Use this widget to select a cluster. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Cluster CPU** and **Cluster Memory**: Use these widgets to view the CPU and memory for the cluster.
- **Cluster IOPS** and **Cluster Network Throughput**: Use these widgets to view the IOPS and network throughput for the cluster.
- Use the other widgets in the dashboard to view which VMs in the cluster generated the highest network throughput and IOPS. You can also view which VMs in the cluster generated the highest CPU demand and the highest memory demand. You can compare the information for the VM with the results for the cluster and correlate the trends. You can manually set the time to the time period for which you want to view data.

Host Utilization Dashboard

The Host Utilization dashboard helps you identify hosts that are extensively consumed from a CPU, memory, disk, and network perspective.

You can use this dashboard to identify hosts that cannot serve the virtual machine demand. The dashboard provides a list of the top 10 virtual machines. You can identify the source of this unexpected demand and take appropriate actions.

You can use the dashboard to view demand patterns over the last 24 hours and identify hosts that have a history of high demand. You must move the virtual machines out of these hosts to avoid potential performance issues. If all the hosts of a given cluster display the same pattern, you might have to add new capacity to cater to the increase in demand.

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 total provisioned capacity, current and historical utilization 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 utilization among vSAN disks from the dashboard. You can view these details either as an aggregate or at an individual cluster level.

Configuration and Compliance Dashboards

The dashboards in the Configuration and Compliance category cater to administrators who are responsible for managing configuration drifts within a virtual infrastructure. Since most of the issues in a virtual infrastructure are a result of inconsistent configurations, dashboards in this category highlight the inconsistencies at various levels such as VMs, hosts, clusters, and virtual networks. You can view a list of configuration improvements that helps you avoid problems that are caused because of misconfigurations.

Your IT security teams can also measure your environment against the vSphere hardening best practices to ensure that your environment is fully secured and meets all the compliance standards.

Key questions these dashboards help you answer are as follows:

- Are the vSphere clusters consistently configured for high availability (HA) and optimal performance?
- Are the ESXi hosts consistently configured and available to use?
- Are the VMs sized and configured as per the recommended best practices?

- Are virtual switches configured optimally?
- Is the environment configured in accordance with the vSphere Hardening Guide?

Cluster Configuration Dashboard

The Cluster Configuration dashboard provides a quick overview of your vSphere cluster configurations. The dashboard highlights the areas that are important in delivering performance and availability to your virtual machines. The dashboard also highlights if there are clusters which are not configured for DRS, High Availability (HA), or admission control to avoid any resource bottlenecks or availability issues when a host fails.

The heat map in this dashboard helps you to identify if you have hosts where vMotion was not enabled as this may not allow the VMs to move from or to that host. This may cause potential performance issues for the VMs on that host if the host gets too busy. You can also view how consistently your clusters are sized and whether the hosts on each of those clusters are consistently configured.

The Cluster Properties widget in this dashboard allows you to report on all these parameters by exporting the data. You can share the data with the relevant stakeholders within your organization.

You can use the dashboard widgets in several ways.

- **vSphere DRS Status**, **vSphere HA Status**, and **HA Admission Control Status**: Use these widgets to view if there are clusters that are not configured for DRS, HA, or admission control. With the information, you can avoid resource bottlenecks or availability issues when a host fails.
- **Is vMotion enabled on hosts in a cluster**: Use this widget to identify if you have hosts where vMotion was not enabled. If vMotion is not enabled, the VMs do not move from or to the host and causes potential performance issues in the VMs on that host if the host gets too busy.
- **Host Count across Clusters**: Use this widget to view all the clusters in your environment. If the clusters have a consistent number of hosts, the boxes displayed are of equal size. This representation helps you determine whether there is a large deviation among cluster sizes, whether there is a small cluster with fewer than four hosts, or whether there is a large cluster. Operationally, keep your clusters consistent and of moderate size.
- **Attributes of ESXi Hosts in the Selected Cluster**: Use this widget to view the configuration details for the hosts within a cluster.
- **All Clusters Properties**: Use this widget to view the properties for all the clusters in the widget.

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.

Network Configuration Dashboard

The Network 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.

VM Configuration Dashboard

The VM dashboard focuses on highlighting the key configurations of the virtual machines in your environment. You can use this dashboard to find inconsistencies in configuration within your virtual machines and take quick remedial measures. You can safeguard the applications which are hosted on these virtual machines by avoiding potential issues due to misconfigurations.

Some of the basic problems the dashboard focuses on includes identifying VMs running on older VMware tools versions, VMware tools not running, or virtual machines running on large disk snapshots. VMs with such symptoms can lead to potential performance issues and hence it is important that you ensure that they do not deviate from the defined standards. This dashboard includes a predefined Virtual Machine Inventory Summary report which you can use to report the configurations highlighted in this dashboard for quick remediation.

You can use the dashboard widgets in several ways.

- Use the Large VMs widgets to view graphical representations of VMs that have a large CPU, RAM, and disk space.
- **Guest OS Distribution**: Use this widget to view a break up of the different flavors of operating systems you are running.

- **Guest Tools Version** and **Guest Tools Status**: Use these widgets to identify if you have inconsistent or older version of VMware tools which might lead to performance issues.
- View the VMs with limits, large snapshots, orphaned VMs, VMs with more than one NIC, and VMs with a nonstandard operating system. These VMs have a performance impact on the rest of the VMs in your environment even though they do not fully use their allocated resources.

You can customize the views in the widgets.

1 Click the **Edit Widget** icon from title bar of the widget. The **Edit** widget dialog box is displayed.

2 From the **Views** section, click the **Edit View** icon. The **Edit View** dialog box is displayed.

3 Click the **Presentation** option in the left pane and make the required modifications.

vSphere Hardening Compliance Dashboard

The vSphere Hardening Compliance dashboard measures your environment against the *vSphere Hardening Guide* and lists any objects which are non-compliant.

This dashboard displays the trend of high risk, medium risk, and low risk violations and shows the overall compliance score of your virtual infrastructure. Using heat maps, you can investigate various components to check the compliance for your ESXi hosts, clusters, port groups, and virtual machines. Each non-compliant object is listed in the dashboard with recommendations on the remediation required to secure your environment.

Operations Dashboards

The dashboards in the Operations category are most helpful to personnel within an organization that require a summary of important data to take quick decisions. As a member of the network operations center (NOC) team, you may want to identify problems and take action or as an executive, you may want a quick overview of your environments to keep track of important KPIs.

Key questions these dashboards help you answer are as follows:

- What does the infrastructure inventory look like?
- What is the alert volume trend in the environment?
- Are virtual machines being served well?
- Are there areas in the data center you have to worry about?
- What does the vSAN environment look like and are there optimization opportunities by migrating VMs to vSAN?

Datastore Usage Overview Dashboard

The Datastore Usage Overview dashboard provides a view of all the virtual machines in your environment in a heat map. The dashboard is suitable for an NOC environment.

The heat map contains a box for each virtual machine in your environment. You can identify the virtual machines that are generating excessive IOPS because the boxes are sized by the number of IOPS they generate.

The colors of the boxes represent the latency experienced by the virtual machines from the underlying storage. An NOC administrator can investigate the cause of this latency and resolve it to avoid potential performance problems.

Host Usage Overview Dashboard

The Host Usage Overview dashboard provides a view of all the ESXi hosts in your environment in a heat map. The dashboard is suitable for an NOC environment.

Using this dashboard an NOC administrator can easily find resource bottlenecks created due to excessive Memory Demand, Memory Consumption or CPU Demand.

The heat map displays hosts grouped by clusters to help you locate clusters that are using excessive CPU or memory. You can also identify if you have ESXi hosts within the clusters that are not evenly utilized. An administrator can then trigger activities such as workload balance or set DRS to ensure that hot spots are eliminated.

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.

Optimize vSAN Deployments

The Optimize vSAN Deployments 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 VMs 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 utilization patterns to see how the VM is served by vSAN after you move the VM.

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.

Performance Troubleshooting Dashboards

The dashboards in the Performance Troubleshooting category cater to the administrators responsible for managing the performance and availability of the virtual machines running in the virtual infrastructure. This category runs you through a guided workflow to answer questions that help you with the troubleshooting process. The dashboards in this category identify and isolate problems that may impact your applications. They provide insight into the full stack to isolate and identify the root cause quickly.

Key questions these dashboards help you answer are as follows:

- Is the application performance impacted due to virtual infrastructure?
- Are noisy neighbors impacting multiple virtual machines and corresponding applications?
- Are there active alerts which require action?
- Are there any known issues impacting the performance and availability of a vSAN cluster?

Troubleshoot a Cluster

The Troubleshoot a Cluster dashboard allows you to identify clusters that have issues and isolate them easily.

You can use the search option to identify a cluster that has an issue. You can also sort the clusters based on the number of active alerts.

After you select the cluster you want to work with, you can view a quick summary of the number of hosts in that cluster and the VMs served by the cluster. The dashboard provides you with current and past utilization trends and also known issues in the cluster in the form of alerts.

You can view the hierarchy of objects related to the cluster and review the status to identify if the objects are impacted because of the current health of the cluster. You can quickly identify any contention issues by looking at the maximum and average contention faced by the VMs on the selected cluster. You can narrow down and view those VMs that have resource contention and take specific steps to troubleshoot and resolve issues.

You can use the dashboard widgets in several ways.

- **Search for a cluster**: Use this widget to select the cluster for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Is your cluster busy?**: Use this widget to view the CPU and memory demand.
- **Are there active alerts on your cluster**: Use this widget to view only the critical alerts.
- **Are the relatives healthy?**: Use this widget to view the hierarchy of the objects related to the cluster and if any of the objects are impacted.
- View the maximum and average CPU, memory, and disk latency for the VMs. If the VM faces contention, it might mean that the underlying infrastructure does not have enough resources to meet the needs of the VMs.
- View a list of VMs that face CPU, memory, and disk latency contention. You can then troubleshoot and take steps to resolve the problem.

Troubleshoot a Datastore

The Troubleshoot a Datastore dashboard allows you to identify storage issues and act on them.

You can use the search option to identify a datastore that has an issue or you can identify a datastore that has high latency as seen in red on the heat map. You can also sort all the datastores with active alerts and troubleshoot the datastore with known issues.

You can select a datastore to see its current capacity and utilization with the number of VMs served by that datastore. The metric charts help you view historical trends of key storage metrics such as latency, outstanding IOs, and throughput.

The dashboard also lists the VMs served by the selected datastore and helps you analyze the utilization and performance trends of those VMs. You can migrate the VMs to other datastores to even out the IO load.

You can use the dashboard widgets in several ways.

- **Search for a datastore**: Use this widget to select the datastore for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the datastore you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Are there active alerts on your datastore**: Use this widget to view only the critical alerts.
- **Are the relatives healthy?**: Use this widget to view the hierarchy of the objects related to the datastore and if any of the objects are impacted.
- **Is your datastore experiencing high latency?** and **Any outstanding disk I/Os?**: Use these widgets to view those datastores with high latency and outstanding disk I/O trends. Ideally, your datastores must not have outstanding disk I/O.
- **How many IOPS is your datastore serving** and **Latency trend for the I/Os done by the VM**: Use these widgets to view the current IOPS and latency of the VMs in the selected datastore.
- Use the other widgets in the dashboard to view trends for the selected datastore regarding disk latency, IOPS, and throughput, VMs served by the datastore and I/O pattern of the selected VM.

Troubleshoot a Host
The Troubleshoot a Host dashboard allows you to search for specific hosts or sort hosts with active alerts. ESXi hosts are the main source of providing resources to a VM and are critical for performance and availability.

To view the key properties of each host, select a host from the dashboard. You can ensure that the host is configured according to the virtual infrastructure design. Any deviation from standards might cause potential issues. You can use the dashboard to answer key questions about current and past utilization and workload trends over the last week. You can also view if the VMs served by the host are healthy.

Since the dashboard lists all the critical events that might affect the availability of the hosts, you can view hardware faults associated with the host. You can view a list of the top 10 VMs that demand CPU and memory resources from the identified host.
Troubleshoot a VM Dashboard
The Troubleshoot a VM dashboard helps an administrator to troubleshoot everyday issues in a virtual infrastructure. While most of the IT issues in an organization are reported at the application layer, you can use the guided workflow in this dashboard to help investigate an ongoing or a suspected issue with the VMs supporting the impacted applications.

You can search for a VM by its name or you can sort the list of VMs with active alerts on them to start your troubleshooting process. When you select a VM, you can view its key properties to ensure that the VM is configured as per your virtual infrastructure design. Any deviation from standards may cause potential issues. You can view known alerts and the workload trend of the VM over the past week. You can also view if any of the resources serving the virtual machine have an ongoing issue.

The next step in the troubleshooting process allows you to eliminate the major symptoms which might impact the performance or availability of a VM. You can use key metrics to find out if the utilization patterns of the VMs are abnormal or if the VM is contending for basic resources such as CPU, memory, or disk.

You can use the dashboard widgets in several ways.

- **Search for a VM**: Use this widget to view all the VMs in the environment. You can select the VM you want to troubleshoot. You can use the filter to narrow your list based on several parameters, such as name, folder name, associated tag, host, or vCenter Server. After you identify the VM you want to troubleshoot, select it. The dashboard is automatically populated with the relevant data.
- **About the VM**: Use this widget to understand the context of the VM. This widget also lends insights to analyze the root cause of the problem or potential mitigations.
- **Are there active alerts on the VM?**: Use this widget to view active alerts. To see noncritical alerts, click the VM object.
- **Is the VM working hard over the last week?**: Use this widget to view the workload trend of the VM for the last week.
- **Are the relatives healthy?**: Use this widget to view the ESXi host where the VM is now running. This host might not be the ESXi host where the VM was running in the past. You can view the remaining related objects and see whether they might contribute to the problem.
- **Is the VMs demand spiking or abnormal?**: Use this widget to identify spikes in the VM demand for any of the resources such as CPU, memory, and network. Spikes in the demand might indicate an abnormal behavior of the VM or that the VM is undersized. The memory utilization is based on the Guest OS metric. It requires VMware Tools 10.0.0 or later and vSphere 6 Update 1 or later. If you do not have these products, the metric remains blank.
- **Is the VM facing contention?**: Use this widget to identify whether the VM is facing contention. If the VM is facing contention, the underlying infrastructure might not have enough resources to meet the needs of the VM.
- **Does the cluster serving the VM have contention?**: Use this widget to view the trend for the maximum CPU contention for a VM within the cluster. The trend might indicate a constant contention within the cluster. If there is contention, you must troubleshoot the cluster as the problem is no longer with the VM.
- **Does the datastore serving the VM have latency?**: Use this widget to help you correlate the latency at the datastore level with the total latency of the VM. If the VM has latency spikes, but the datastore does not have such spikes, it might indicate a problem with the VM. If the datastore faces latency as well, you can troubleshoot to find out why the datastore has these spikes.
- **Parent Host** and **Parent Cluster**: Use these widgets to view the host and the cluster on which the VM resides.

Troubleshoot vSAN Dashboard

The Troubleshoot vSAN dashboard helps you view the properties of your vSAN cluster and the active alerts on the cluster components. The cluster components include hosts, disk groups, or the vSAN datastores.

You can select a cluster from the dashboard and then list all the known problems with the objects associated with the cluster. The objects include clusters, datastores, disk groups, physical disks, and VMs served by the selected vSAN cluster.

You can view the key utilization 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.
- **Is the cluster busy?**: Use this widget to view the basic utilization of the cluster in terms of CPU, memory, capacity, components, throughput, and IOPS. You can also view the past trend to determine the normal usage or high utilization period.
- **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 cache tier congested?**: Use this widget to view if the cache tier is congested in a cluster.
- **Is the write buffer full on diskgroups?**: Use this widget to view the usage of the write buffer on diskgroups in a cluster.
- **Are the read IOs served by cache?**: Use this widget to determine if the read IOs are served by cache. This is not applicable to an all flash cluster type.
- **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.

- **Is the selected host busy?**: Use this widget to view CPU and memory workload, CPU and memory contention, network usage, and number of packets dropped.
- **Cache Disks: Any hardware issues?**: Use this widget to view the individual cache disks measured against various metrics.
- **Capacity Disks: Any hardware issues?**: Use this widget to view the individual capacity disks measured against various metrics.

Troubleshoot with Logs Dashboard
When vRealize Operations Manager is integrated with vRealize Log Insight, you can access the custom dashboards and content pack dashboards from the Troubleshoot with Logs dashboard. You can view graphs of log events in your environment, or create custom sets of widgets to access the information that matters most to you.

You can investigate an ongoing issue within your virtual infrastructure using the logs. You can view predefined views created within vRealize Log Insight to answer questions from predefined queries within vRealize Log Insight.

You can correlate metrics and queries within vRealize Operations Manager to troubleshoot issues across applications and infrastructure.

For more information about the Troubleshoot with Logs dashboard, see the vRealize Log Insight documentation.

To access the Troubleshoot with Logs dashboard from vRealize Operations Manager, you must either:

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

  For more information on configuring, see Configuring vRealize Log Insight with vRealize Operations Manager.

vRealize Automation Dashboards

With the vRealize Automation dashboards, you can monitor and troubleshoot objects in your cloud infrastructure.

The following vRealize Automation solution dashboards are added to the predefined vRealize Operations Manager dashboards:

- vRealize Automation Environment Overview
- vRealize Automation Top-N

vRealize Automation Environment Overview Dashboard
You can use the vRealize Automation Environment Overview dashboard to view information about the tenants and the related alerts. You can also view the health, risk, and efficiency of resources for a given object from the managed inventory.

You can use the vRealize Automation Tenant 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.

- **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.
- **Environment Overview**. 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.
- **Blueprint List**. Use this widget to view the blueprint 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.
- **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.
- **vRealize Automation Managed Cluster**. 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.

vRealize Automation Top-N Dashboard

You can use the widgets in the vRealize Automation Top-N dashboard to view the top results from analysis of blueprints and business groups that you select.

You can use the vRealize Automation Top-N dashboard to perform some of the following tasks:

- To view the most popular blueprints.
- To view the business groups that have the most critical alerts.

You can use the dashboard widgets in several ways.

- **Business Groups with most Critical Alerts**. Use this widget to view the business groups that have the most critical alerts.
- **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 over a seven-day period. You can create one or more graphs that contain metrics for the object you select.
- **Most Popular Deployed Blueprints**. Use this widget to view the top five most popular deployed blueprints in the environment.

Create and Configure Dashboards

To view the status of all objects in vRealize Operations Manager, create a dashboard by adding widgets. You can create and modify dashboards and configure them to meet your environment needs.

Procedure

1. In the menu, click **Dashboards**.
2. Click **Actions > Create Dashboard** to create and configure a dashboard.
3. Complete the steps in the left pane to:
   a. Enter a name for the dashboard.

      Name and Description Details

   b. Add widgets to the dashboard.

      Widget List Details

   c. Configure widget interactions.

      Widget Interactions Details

   d. Create dashboard navigation.

      Dashboard Navigation Details
4. Click **Save**.
5. Click **Actions > Edit Dashboard** to modify the dashboard.

Name and Description Details

The name and visualization of the dashboard as it appears on the vRealize Operations Manager Home page.

Where You Configure a Dashboard

To create or edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. In the workspace, on the left, click **Dashboard Configuration**.

Table 4-137. Dashboard Configuration Options in the Dashboard Workspace

Option	Description
Name	Name of the dashboard as it appears on top of the tab on the Home page and in the dashboard's lists. 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`.
Description	Description of the dashboard.
Is default	If you select **Yes**, the dashboard appears on the Home page when you log in.

Widget List Details

vRealize Operations Manager provides a list of widgets that you can add to your dashboard to monitor specific metrics and properties of objects in your environment.

Where You Add Widgets to a Dashboard

To create or edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. In the workspace, on the left, click **Widget List**. If you create a dashboard, complete the required previous steps of the workspace.

How to Add Widgets to a Dashboard

In the workspace, on the left, you see a list with all the predefined vRealize Operations Manager widgets. To add a widget to the dashboard, drag the widget to the content area on the right.

To locate a widget, you can type the name or part of the name of a widget in the **Filter** option. For example, when you enter **`cap`**, the list is filtered to display the Capacity Remaining, Capacity Utilisation, and Reclaimable Capacity widgets. You can then select the widget you require.

Most widgets must be configured individually to display information. For more information about how to configure each widget, see Widgets.

How to Arrange Widgets in a Dashboard

You can modify your dashboard layout to suit your needs. By default, the first widgets that you add are automatically arranged horizontally wherever you place them. The widgets move up to the highest position in the dashboard based on their width.

- To position a widget, drag the widget to the desired location in the layout. Other widgets automatically rearrange to make room.
- To resize a widget, drag the bottom right corner of the widget.

Widget Interactions Details

You can connect widgets so that the information they show depends on each other.

Where You Create Widget Interactions

To create a widget interaction for widgets in a dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. In the workspace, on the left, click **Widget Interactions**. If you create a new dashboard, complete the required previous steps of the workspace.

How to Create Widget Interactions

The list of available widget interactions depends on the widgets in the dashboard. Widgets can provide, receive, and do both. Some widgets can have more than one provider.

To create interactions, click the **Selected Object(s)** drop-down menu for the specified widget and select the provider widget. There are widgets that provide alerts, metrics, or tags. Click the **Selected Alert(s)**, **Selected Metric(s)**, or **Selected Tag(s)** drop-down menu to select the alert, metric, or tag specific provider widget. When you are ready with all interactions, click **Apply Interactions**. For more information about how interactions work, see Widget Interactions.

Dashboard Navigation Details

You can use dashboard navigation to move from one dashboard to another, and to apply sections or context from one dashboard to another. You can connect a widget to widgets on other dashboards to investigate problems or better analyze the provided information.

Where You Add Dashboard Navigation

To create a dashboard navigation to a dashboard, in the menu, click **Dashboards**. Click **Actions > Create Dashboard** to add a dashboard or **Actions > Edit Dashboard** to edit the selected dashboard. In the workspace, on the left, click **Dashboard Navigation**. If you create a dashboard, complete the required previous steps of the workspace.

How Dashboard Navigation Works

You can create dashboard navigation only for provider widgets. The provider widget sends information to the destination widget. When you create dashboard navigation, the destination widgets are filtered based on the information type they can receive.

How to Add a Dashboard Navigation to a Dashboard

The list of available dashboard navigation depends on the available dashboards and the widgets in the current dashboard. To add navigation, click the **Destination Dashboards** drop-down menu for the specified widget and select the dashboard and the widget to navigate to. You can select more than one applicable widget. Click **Apply Navigations** to apply the connections.

Note If a dashboard is unavailable at the Home page, it is unavailable for dashboard navigation.

The Dashboard Navigation icon ( ) appears in the top menu of each widget when a dashboard navigation is available. You can select multiple objects to apply selections or context from one dashboard to another. Press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.

Managing Dashboards

You can change the order of the dashboard tabs, configure vRealize Operations Manager to switch from one dashboard to another, create dashboard folders to group the dashboards in a way that is meaningful to you, and share a dashboard or dashboard template with one or more user groups.

Reorder and Switch Dashboards

You can change the order of the dashboard tabs on your home page. You can configure vRealize Operations Manager to switch from one dashboard to another. This feature is useful if you have several dashboards that show different aspects of your enterprise's performance and you want to look at each dashboard in turn.

Where You Configure a Dashboard Order and Automatic Switch

To reorder and configure a dashboard switch, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Reorder/Autoswitch Dashboards**.

How You Reorder the Dashboards

The list shows the dashboards as they are ordered. Drag the dashboards up and down to change their order on the home page.

How You Configure an Automatic Dashboard Switch

1 Double-click a dashboard from the list to configure.

2 From the Auto Transition drop-down menus, select **On**.

3 Select the switch time interval in seconds.

4 Select the dashboard to switch to and click **Update**.

5 Click **Save** to save your changes.

On the home page, the current dashboard will switch to the dashboard that is defined after the specified time interval.

Manage Summary Dashboards

The **Summary** tab provides you with an overview of the state of the selected object, group, or application. You can change the **Summary** tab with a dashboard to get information specific to your needs.

Where You Configure a Summary Tab Dashboard

To manage the summary dashboards, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Manage Summary Dashboards**.

How You Manage the Summary Tab Dashboard

Table 4-138. Manage Summary Dashboards Options

Option	Description
Adapter Type	Adapter type for which you configure a summary dashboard.
Filter	Use a word search to limit the number of adapter types that appear in the list.
Name	List with all available objects.
Use Default icon	Click to use vRealize Operations Manager default **Summary** tab.
Detail Page	Shows what kind of **Summary** tab you use for the selected object.
Assign a Dashboard icon	Click to view the Dashboard List dialog box that lists all the available dashboards.

To change the Summary tab for an object, select the object in the left panel, click the **Assign a Dashboard** icon. Select a dashboard for it from the Dashboard List dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard you have associated to the object type when you navigate to the **Summary** tab of the object details page.

Manage Dashboard Groups

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

Where You Configure a Dashboard Group

To manage the dashboard groups, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Manage Dashboard Groups**.

How You Manage the Dashboard Groups

Table 4-139. Manage Tap Groups Options

Option	Description
Dashboard Groups	A hierarchy tree with all available group folders.
Dashboards List	A list with all available dashboards.

To create a dashboard group folder, right-click the **Dashboard Groups** folder or another folder and click **Add**. To add a dashboard, drag one from the Dashboards list to the folder.

Share Dashboards

You can share a dashboard or dashboard template with one or more user groups. When you share a dashboard, it becomes available to all of the users in the user group that you select. The dashboard appears the same to all of the users who share it. If you edit a shared dashboard, the dashboard changes for all users. Other users can only view a shared dashboard. They cannot change it.

Where You Share a Dashboard From

To share a dashboard, in the menu, click **Dashboards**. Select **Actions > Manage Dashboards**. Click the gear icon and select **Share Dashboards**.

Table 4-140. Share Dashboards Options

Option	Description
Accounts Group	All available groups with which you can share a dashboard.
Shared Dashboards	All available dashboards and templates that you can share. You can switch between dashboard tabs and dashboard templates by clicking the **Share Dashboard Tabs/ Templates** icon.

How You Manage a Shared Dashboard Tab

To share a dashboard tab, navigate to the dashboard in the list of Shared Dashboards and drag it to the group to share it with on the left.

To stop sharing a dashboard with a group, click that group on the left panel, navigate to the dashboard in the right panel, and click the **Stop Sharing** icon above the list.

To stop sharing a dashboard with more than one group, click the **Not Grouped** name on the left panel, navigate to the dashboard in the right panel, and click the **Stop Sharing** icon above the list.

Views

vRealize Operations Manager provides several types of views. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective. vRealize Operations Manager Views also show information that the adapters in your environment provide.

You can configure vRealize Operations Manager views to show transformation, forecast, and trend calculations.

- The transformation type determines how the values are aggregated.
- The trend option shows how the values tend to change, based on the historical, raw data. The trend calculations depend on the transformation type and roll up interval.
- The forecast option shows what the future values can be, based on the trend calculations of the historical data.

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.

In the menu, click **Dashboards**, and then in the left pane click **Views** to access the Views page.

On the Views page you can create, edit, delete, clone, export, and import views.

You can order the listed views by name, type, description, subject, or owner.

You can limit the views list by adding a filter from the upper-right corner of the panel.

Table 4-141. Filter Groups

Filter Group	Description
Name	Filter by the view name. For example, type **my view** to list all views that contain the **my view** phrase in their name.
Type	Filter by the view type.
Description	Filter by the view description. For example, type **my view** to list all views that contain the **my view** phrase in their description.
Subject	Filter by the subject.
Owner	Filter by the owner.

Views and Reports Ownership

Views, reports, templates, or schedules owner might change in time.

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Create and Configure a View

To collect and display information for a specific object, you can create a custom view.

Procedure

1. In the menu, click **Dashboards**, and then in the left pane click **Views**.
2. Click the **Create View** icon to create a view.
3. Complete the steps in the left pane to:
   a. Enter a name and description for the view.

      Name and Description Details

   b. Change the presentation of a view.

      Presentation Details

   c. Select the base object type for a view.

      Subjects Details

   d. Add data to a view.

      Data Details

   e. Change the visibility of a view.

      Visibility Details
4. Click **Save**.
5. From the Views page, click the **Edit View** icon to modify the view.

Name and Description Details

The name and description of the view as they appear in the list of views on the Views page.

To add a name and description to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Name and Description**.

Table 4-142. Name and Description Options in the View Workspace

Option	Description
Name	Name of the view as it appears on the Views page.
Description	Description of the view.

Presentation Details

A presentation is a way the collected information for the object is presented. Each type of view helps you to interpret metrics and properties from a different perspective.

To change the presentation of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Presentation**. If you create a view, complete the required previous steps.

Table 4-143. Presentation Options in the View Workspace

View Type	Description
List	Provides tabular data about specific objects in the monitored environment.
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.
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 4-144. 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.
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.

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 4-145. Dynamic Distribution Configuration Options

Configuration Option	Description
Buckets Count	The number of buckets to use in the data distribution.
Buckets Size Interval	The bucket size is determined by the defined interval divided by the specified number of buckets.

Table 4-145. Dynamic Distribution Configuration Options (continued)

Configuration Option	Description
Buckets Size Logarithmic bucketing	The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data.
Buckets Size Simple Max/Min bucketing	The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets.

Manual distribution

You specify the number of buckets and the minimum and maximum values of each bucket.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distribute the data.

View Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Dynamic distribution

You specify in details how vRealize Operations Manager distributes the data in the buckets.

Table 4-146. Dynamic Distribution Configuration Options

Configuration Option	Description
Buckets Count	The number of buckets to use in the data distribution.
Buckets Size Interval	The bucket size is determined by the defined interval divided by the specified number of buckets.
Buckets Size Logarithmic bucketing	The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data.
Buckets Size Simple Max/Min bucketing	The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets.

Manual distribution

You specify the number of buckets and the minimum and maximum values of each bucket.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distribute the data.

If you increase the number of buckets, you can see more detailed data.

Subjects Details

The subject is the base object type for which the view shows information.

To specify a subject for a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Subjects**. If you create a new view, complete the required previous steps.

The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. You can limit the level where the view appears with the Blacklist option in the **Visibility** step.

View availability depends on the view configuration subject, inventory view, user permissions, and view Visibility settings.

For list views with **Symptom** as a subject, the following columns can be sorted: Criticality Level, Status, Object Type, Object Name, Created on, and Canceled on. You cannot sort the Triggered On and Violation Info columns. If other symptom metrics exist, you cannot sort any of the columns.

Views Applicability

List View

When you navigate through the environment tree, you can see the List view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the List view might be missing at the object containers. For example, you create a List view with subject Host System. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Hosts and Clusters > vSphere World**. Select a vCenter Server, and click the **Details** tab, you can see your List view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Storage > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your List view is missing. Your List view with subject Host System is missing because the object Host System is not included in the vSphere Storage inventory view.

Summary View

When you navigate through the environment tree, you can see the Summary view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the Summary view might be missing at the object containers. For example, you create a Summary view with subject Datastore. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Storage > vSphere World**. Select a vCenter Server, and click the **Details** tab, you can see your List view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Networking > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your Summary view is missing. Your Summary view with subject datastore is missing because the object Datastore is not included in the vSphere Networking inventory view.

Trend View

When you navigate through the environment tree, you can see the Trend view only at the subjects that you specify during the view configuration. For example, you create a Trend view with subject Virtual Machine. When you navigate to a virtual machine in the navigation tree, you see your view.

Distribution View

When you navigate through the environment tree, you can see the Distribution view only at the object containers of the subjects that you specify during the view configuration. Depending on the inventory view, the Distribution view might be missing at the object containers. For example, you create a Distribution view with subject Host System. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Hosts and Clusters > vSphere World**. Select a vCenter Server, and click the **Details** tab, you can see your Distribution view. In the menu, click **Environment**, and then in the left pane click **vSphere Environment > vSphere Networking > vSphere World**. Select the same vCenter Server, and click the **Details** tab, your Distribution view is missing. Your Distribution view with subject Host System is missing because the object Host System is not included in the vSphere Networking inventory view.

Text View

When you navigate through the environment tree, you can see the Text view only at the subjects that you specify during the view configuration. For example, you create a Text view with subject vCenter Server. When you navigate to a vCenter Server in the navigation tree, you see your view. If you did not specify a subject, you see your view for every subject in the environment.

Image View

The Image view is applicable for every object in the environment.

Note Views applicability depends also on your user permissions and the view Visibility configuration.

Views Applicability

Views might not always appear where you expect them to. The main applicability of views depends on the view subject and the inventory view.

List View

When you navigate through the environment tree, you can see the List view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the List view might be missing at the object containers. For example, you create a List view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Storage > vSphere World**, select the same vCenter Server, and click the **Details** tab, your List view is missing. Your List view with

subject Host System is missing because the object Host System is not included in the vSphere Storage inventory view.

Summary View

When you navigate through the environment tree, you can see the Summary view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the Summary view might be missing at the object containers. For example, you create a Summary view with subject Datastore. When you go to **Environment > vSphere Storage > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Summary view is missing. Your Summary view with subject datastore is missing because the object Datastore is not included in the vSphere Networking inventory view.

Trend View

When you navigate through the environment tree, you can see the Trend view only at the subjects that you specify during the view configuration. For example, you create a Trend view with subject Virtual Machine. When you navigate to a virtual machine in the navigation tree, you see your view.

Distribution View

When you navigate through the environment tree, you can see the Distribution view only at the object containers of the subjects that you specify during the view configuration. Depending on the inventory view, the Distribution view might be missing at the object containers. For example, you create a Distribution view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your Distribution view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Distribution view is missing. Your Distribution view with subject Host System is missing because the object Host System is not included in the vSphere Networking inventory view.

Text View

When you navigate through the environment tree, you can see the Text view only at the subjects that you specify during the view configuration. For example, you create a Text view with subject vCenter Server. When you navigate to a vCenter Server in the navigation tree, you see your view. If you did not specify a subject, you see your view for every subject in the environment.

Image View

The Image view is applicable for every object in the environment.

Note Views applicability depends also on your user permissions and the view Visibility configuration.

Data Details

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which vRealize Operations Manager collects, calculates, and presents the information for the view.

To add data to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data**. If you create a new view, complete the required previous steps.

How to Add Data to a View

If you selected more than one subject, specify the subject for which you add data. Double-click the data from the tree in the left panel to add it to the view. For each subject the data available to add might be different.

How to Configure the Data Transformation

The data configuration options depend on the view and data type that you select. Most of the options are available for all views.

Table 4-147. Data Configuration Options

Configuration Option	Description
Metric name	Default metric name. Available for all views.
Metric label	Customizable label as it appears in the view or report. Available for all views.
Units	Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU| Demand(MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. Available for all views.
Sort order	Orders the values in ascending or descending order. Available for List view and Summary view.

Table 4-147. 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. ■ Last. Ignores all the data except the data you receive most recently and that is within the selected time range. ■ 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. Available for all views, except Trend.
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.
Projects	A project contains scenarios and is a supposition about how capacity and load change if certain conditions are changed without making actual changes to your virtual infrastructure. If you implement the project, you know in advance what your capacity requirements are. Available for all views. Depends on the selected metrics and properties.

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 4-148. 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 4-149. 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 will show 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 new view, complete the required previous steps.

Each subject has a separate filter box. For Alerts Rollup, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 4-150. Filter Add Options

Option	Description
Add	Adds another criteria to the criteria set. The filter returns results that match all of the specified criteria.
Add another criteria	Adds another criteria set. The filter returns results that match one criteria set or another.

How to Add a Summary Row or Column to a View

The summary option is available only for List and Summary views. It is mandatory for the Summary views. You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, select **Content > Views** in the left pane. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data** and click the **Summary** tab in the main panel. If you create a new view, complete the required previous steps.

For the List view, the summary row shows aggregated information by the specified subjects.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

Visibility Details

The view visibility defines where you can see a view in vRealize Operations Manager.

To change the visibility of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Visibility**. If you create a new view, complete the required previous steps.

Table 4-151. 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.
Further Analysis	Select a badge to make the view available at Further Analysis. Further Analysis section appears on the Analysis tab of an object. When you make a view visible for a badge, a link to it appears in the Further Analysis section of that badge. You can click on the link to analyse the provided information.
Blacklist	Select a subject level where you do not want to see this view. For example, you have a list view with subject virtual machines. It is visible when you select any of its parent objects. You add datacenter in the blacklist. The view is not visible anymore on datacenter level.

Editing, Cloning, and Deleting a View

You can edit, clone, and delete a view. Before you do, familiarize yourself with the consequences of these actions.

When you edit a view, all changes are applied to the report templates that contain it.

When you clone a view, the changes that you make to the clone do not affect the source view.

When you delete a view, it is removed from all the report templates that contain it.

User Scenario: Create, Run, Export, and Import a vRealize Operations Manager View for Tracking Virtual Machines

As a virtual infrastructure administrator, you use vRealize Operations Manager to monitor several environments. You must know the number of virtual machines on each vCenter Server instance. You define a view to gather the information in a specific order and use it on all vRealize Operations Manager environments.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

You will create a distribution view and run it on the main vRealize Operations Manager environment. You will export the view and import it in another vRealize Operations Manager instance.

Procedure

1 Create a vRealize Operations Manager View for Supervising Virtual Machines

   To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

2 Run a vRealize Operations Manager View

   To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

3 Export a vRealize Operations Manager View

   To use a view in another vRealize Operations Manager, you export a content definition XML file.

4 Import a vRealize Operations Manager View

   To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Create a vRealize Operations Manager View for Supervising Virtual Machines

To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

Procedure

1 In the menu, click **Dashboards**, and then in the left pane click **Views**.

2 Click the plus sign to create a new view.

3 Enter **`Virtual Machines Distribution`**, the name for the view.

4 Enter a meaningful description for the view.

For example, **`A view showing the distribution of virtual machines per hosts.`**

5 Click **Presentation** and select the **Distribution** view type.

The view type is the way the information is displayed.

a From the **Visualization** drop-down menu, select **Pie Chart**.

b From the Distribution Type configurations, select **Discrete distribution**.

Leave **Max number of buckets** deselected because you do not know the number of hosts on each vCenter Server instance. If you specify a number of buckets and the hosts are more than that number, one of the slices shows unspecified information labeled Others.

6 Click **Subjects** to select the object type that applies to the view.

a From the drop-down menu, select **Host System**.

The Distribution view is visible at the object containers of the subjects that you specify during the view configuration.

7 Click **Data** and in the filter text box enter **`Total Number of VMs`**.

8 Select **Summary > Total Number of VMs** and double-click to add the metric.

9 Retain the default metric configurations and click **Save**.

Run a vRealize Operations Manager View

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

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

2 In the left pane, navigate to a vCenter Server instance and click the **Details** tab.

All listed views are applicable for the vCenter Server instance.

3 From the **All Filters** drop-down menu on the left, select **Type > Distribution**.

You filter the views list to show only distribution type views.

4 Navigate to and click the **Virtual Machines Distribution** view.

The bottom pane shows the distribution view with information about this vCenter Server. Each slice represents a host and the numbers on the far left show the number of virtual machines.

Export a vRealize Operations Manager View

To use a view in another vRealize Operations Manager, you export a content definition XML file.

If the exported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

1. In the menu, click **Dashboards**, and then in the left pane click **Views**.
2. Click the gear icon and select **Export View**.
3. In the list of views, navigate to and click the **Virtual Machines Distribution** view .
4. Select **All Actions > Export view**.
5. Select a location on your local system to save the XML file and click **Save**.

Import a vRealize Operations Manager View

To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

1. In the menu, click **Dashboards**, and then in the left pane click **Views**.
2. Click the gear icon and select **Import View**.
3. Browse to select the Virtual Machines Distribution content definition XML file and click **Import**.

   If the imported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

   Note The imported view overwrites if a view with the same name exists. All report templates that use the existing view are updated with the imported view.

Reports

A report is a scheduled snapshot of views and dashboards. You can create it to represent objects and metrics. It can contain table of contents, cover page, and footer.

With the vRealize Operations Manager reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.

Report Templates Tab

On the **Report Templates** tab you can create, edit, delete, clone, run, schedule, export, and import templates.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Report Templates** to access the Reports Templates tab.

All templates that are applicable for the selected object are listed on the **Report Templates** tab. You can order them by report name, subject, date they were modified, last run, or owner.

You can filter the templates list by adding a filter from the right side of the panel.

Table 4-152. 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.

vSphere users must be logged in until the report generation is complete. If you log out or your session expires, the report generation fails.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

Generated Reports Tab

All reports that are generated for a selected object are listed on the **Generated Reports** tab.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Generated Reports** to access the Generated Reports tab.

You can order the reports by the date and time that they were created, the report name, the owner, or their status. If the report is generated through a schedule, the owner is the user who created the schedule.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

You can filter the reports list by adding a filter from the right side of the panel.

Table 4-153. 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.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by those objects.

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

Create and Modify a Report Template

You create a report to generate a scheduled snapshot of views and dashboards. You can track current resources and predict potential risks to the environment. You can schedule automated reports at regular intervals.

Procedure

1. In the menu, click **Dashboards**, and then in the left pane click **Reports**.
2. On the **Report Templates** tab, click the **New Template** icon to create a template.
3. Complete the steps in the left pane to:
   a. Enter a name and description for the report template.

      Name and Description Details

   b. Add a view or a dashboard.

      Views and Dashboards Details

   c. Select an output for the report.

      Formats Details

   d. Select the layout options.

      Layout Options Details
4. Click **Save**.
5. From the Report Templates tab, click **Edit Template** to modify the report template.

Name and Description Details

The name and description of the report template as they appear in the list of templates on the **Report Templates** tab.

Where You Add Name and Description

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Name and Description**.

Table 4-154. Name and Description Options in the Report Template Workspace

Option	Description
Name	Name of the template as it appears on the **Report Templates** tab.
Description	Description of the template.

Views and Dashboards Details

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs.

Where You Add Views and Dashboards

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Views and Dashboards**. If you create a template, complete the required previous steps of the workspace.

How You Add Views and Dashboards

To add a view or a dashboard to your report template, select it from the list on the left pane and drag it to the main panel. You can drag the views and dashboards in the main panel to reorder them. You can select a portrait or landscape orientation for each view or dashboard from the drop-down menu next to its title.

Table 4-155. 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.

Table 4-155. Views and Dashboards Options in the Report Template Workspace (continued)

Option	Description
Create Dashboard	Create a dashboard directly from the template workspace. This option is available when you select **Dashboards** from the **Data type** drop-down menu.
Edit Dashboard	Edit a dashboard directly from the template workspace. This option is available when you select **Dashboards** from the **Data type** drop-down menu.
Search	Search for views or dashboards by name. To see the complete list of views or dashboards, delete the search box contents and press Enter.
List of views	List of the views that you can add to the template. This list is available when you select **Views** from the **Data type** drop-down menu.
List of dashboards	List of the dashboards that you can add to the template. This list is available when you select **Dashboards** from the **Data type** drop-down menu.
Preview of views and dashboards	In the main panel, you see a preview of the views and dashboards that you add. When you create a template in the context of an object from the environment, you see a live preview of the views and dashboards.

Formats Details

The formats are the outputs in which you can generate the report.

Where You Add Formats

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Formats** to select a format for the report template. If you create a template, complete the required previous steps of the workspace.

Table 4-156. Formats Options in the Report Template Workspace

Option	Description
PDF	With the PDF format, you can read the reports, either on or off line. This format provides a page-by-page view of the reports, as they appear in printed form.
CSV	In the CSV format, the data is in a structured table of lists.

Layout Options Details

The report template can contain layout options such as a cover page, table of contents, and footer.

Where You Add Layout Options

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. On the Report Templates toolbar, click the **New Template** icon to add a template or the **Edit Template** icon to edit the selected template. From the New Template or Edit Report Template dialog box, in the workspace, on the left, click **Layout Options**. If you create a template, complete the required previous steps of the template.

Table 4-157. Layout Options in the Report Template Workspace

Option	Description
Cover Page	Can contain an image up to 5 MB. The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page.
Table of contents	Provides a list of the template parts, organized in the order of their appearance in the report.
Footer	Includes the date when the report is created, a note that the report is created by VMware vRealize Operations Manager, and page number.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

1. In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.
2. From the toolbar, click the **Add** icon.
3. From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

   The dialog box expands to include your plug-in instance settings.
4. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.
5. Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.

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 > Report Templates** to access the Report Templates tab.

On the **Report Templates** tab, you can create, edit, delete, clone, run, schedule, export, and import templates.

The listed templates are user-defined and predefined by vRealize Operations Manager. You can order them by template name, subject, date they were modified, last run, or owner. For each template, you can see the number of generated reports and schedules.

You can filter the templates list by adding a filter from the right side of the panel.

Table 4-158. 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.

The maximum number of reports per template is 10. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

Generated Reports Overview

A report is a scheduled snapshot of views and dashboards. It presents data in downloadable formats.

In the menu, click **Dashboards**, and then in the left pane select **Reports > Generated Reports** to access the Generated Reports tab.

The list contains all generated reports. You can order them by the date and time they were created, report name, owner, or status. If the report is generated through a schedule, the owner is the user who created the schedule.

Note The maximum number of reports per template is 10. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

You can filter the reports list by adding a filter from the upper-right corner of the panel.

Table 4-159. 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.

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, and click the **gear icon** > **Schedule report**.

How Do You Schedule a Report

Table 4-160. 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 an FTP server. For more information about how to set up and configure the email options, see Outbound Settings. 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 subfolder 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.

User Scenario: Handling Reports to Monitor Virtual Machines

As a virtual infrastructure administrator, you use vRealize Operations Manager to monitor several environments. You must present to your team a report with your corporate logo for all oversized and stressed virtual machines, and their current and trend memory use. You use predefined report templates to gather and format the information in a specific order.

You will create a report template with predefined views and dashboards. You will generate the report to test the template and create a schedule for generating the report once every two weeks.

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 Create a Report Template for Monitoring Virtual Machines

To monitor oversized and stressed virtual machines, and their memory use, you create a report template.

2 Generate a Report

To generate a report, you use the Virtual Machines Report template for a vCenter Server system that shows information for oversized and stressed virtual machines, and their memory use.

3 Download a Report

To verify that the information appears as expected, you download the generated report from the Virtual Machines Report template.

4 Schedule a Report

To generate a report on a selected date, time, and recurrence, you create a schedule for the Virtual Machines Report template. You set the email options to send the generated report to your team.

Create a Report Template for Monitoring Virtual Machines

To monitor oversized and stressed virtual machines, and their memory use, you create a report template.

You create a report template with PDF and CSV output and add views, dashboards, and layout options to it.

Prerequisites

- Understand the concept of vRealize Operations Manager views. See Views.
- Know the location of your corporate logo.

Procedure

Procedure

1 In the menu, click **Dashboards**, and then in the left pane click **Reports**.

2 On the **Report Templates** tab, click the **New Template** icon to create a template.

3 Enter **`Virtual Machines Report`**, the name for the template.

4 Enter a meaningful description for the template.

For example,
`A template for oversized and stressed virtual machines, and their memory use.`

5 Click **Views and Dashboards**. On the **Data type** drop-down menu, leave **Views** selected.

The currently configured views are available in the list below the **Data type** drop-down menu. Views present collected information for an object in a certain way depending on the view type.

6 In the search box, enter `Virtual Machine`.

The list is now limited to views where the name contains Virtual Machine.

7 Double-click the views to add them to the template.

Option	Description
Virtual Machine Rightsizing CPU, Memory, and Disk Space	Monitors oversized VMs
Virtual Machine Recommended CPU and Memory Size	Monitors stressed VMs

The views appear in the main panel of the workspace with a preview of sample data.

8 In the search box, enter `VM`.

The list is now limited to views where the name contains VM.

9 Navigate to *VMs Memory Usage (%) Distribution* view, and double-click the view to add it to the template.

The view appears in the main panel of the workspace with a preview of sample data.

10 (Optional) In the main panel of the workspace, drag the views up and down to reorder them.

11 From the **Data type** drop-down menu, select **Dashboards**.

The currently configured dashboards appear in the list below the **Data type** drop-down menu. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure.

12 Double-click **vSphere VMs Memory**, **vSphere VMs CPU**, and **vSphere VMs Disk and Network** dashboards to add them to the template.

The dashboards appear in the main panel of the workspace.

13 Click **Formats** and leave the **PDF** and **CSV** check boxes selected.

14 Click **Layout Options** and select the **Cover Page** and **Footer** check boxes.

The corresponding panes appear in the main panel of the workspace.

15 In the Cover Page panel, click **Browse** and navigate to an image on your computer.

The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page.

The image uploads to a database. It is used for the cover page every time you generate a report from this template.

16 Click **Save**.

Results

Your report template is saved and listed on the **Report Templates** tab.

What to do next

Generate and download the report to verify the output. See Generate a Report

Generate a Report

To generate a report, you use the Virtual Machines Report template for a vCenter Server system that shows information for oversized and stressed virtual machines, and their memory use.

Prerequisites

Create a report template. See Create a Report Template for Monitoring Virtual Machines.

Procedure

1. In the menu, click **Environment**.
2. In the left pane, navigate to a vCenter Server system.
3. Click the **Reports** tab and click **Report Templates**.

   The listed report templates are associated with the current object.
4. Navigate to the **Virtual Machines Report** template and click the **Run Template** icon.

Results

The report is generated and listed on the **Generated Reports** tab.

What to do next

Download the generated report and verify the output. See Download a Report.

Download a Report

To verify that the information appears as expected, you download the generated report from the Virtual Machines Report template.

Prerequisites

Generate a report from the Virtual Machines Report template. See 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 ( ) and CSV ( ) icon to save the report in the relevant file format.

Results

vRealize Operations Manager saves the report file to the location you selected.

What to do next

Schedule a report generation and set the email options, so your team receives the report. See Schedule a Report.

Schedule a Report

To generate a report on a selected date, time, and recurrence, you create a schedule for the Virtual Machines 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 vRealize Operations Manager 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 vCenter Server .
3. Click the **Reports** tab and click **Report Templates**.
4. Select the **Virtual Machines Report** template from the list.
5. Click the gear icon ( ) and select **Schedule report**.
6. Select the time zone, date, and hour 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 **Weekly** and set the report generation for every two weeks on Monday.
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.
   b. Select an outbound rule.

   An email is sent according to this schedule every time a report is generated.
9. 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.

Configuring Administration Settings

After vRealize Operations Manager is installed and configured, you can use administration settings to manage your environment. You find most administration settings under the Administration selection of the vRealize Operations Manager interface.

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 may monitor and when your licenses expire.

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.

Use the toolbar options to add, edit, or remove items.

Table 4-161. License Key Toolbar Options

Option	Description
Add	Select a solution or product, and enter and validate a license key for it.
Delete	Remove a license key.
Refresh	Update the list of keys.

Use the data grid options to view item details.

Table 4-162. 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
License Capacity	Number of objects that the license allows the product to monitor
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 currently 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 Group Options

The options include toolbar and data grid options.

Use the toolbar options to add, edit, or remove items.

Table 4-163. 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 4-164. 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** .

Use the toolbar options to add, edit, or remove items.

Table 4-165. Maintenance Schedule Toolbar Options

Option	Description
Add	Open a window in which you can select the maintenance schedule settings for a new schedule.
Edit	Open a window in which you can 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 4-166. 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

Table 4-166. Manage Maintenance Schedule Add or Edit Options (continued)

Option	Description
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.
- Use VMware vCenter Server users. After the vCenter Server is registered with vRealize Operations Manager, configure the vCenter Server user options in the vRealize Operations Manager global settings to enable a vCenter Server user to log in to vRealize Operations Manager. When logged into vRealize Operations Manager, vCenter Server users access objects according to their vCenter Server-assigned permissions.
- Add an authentication source to authenticate imported users and user group information that resides on another machine.
  - Use LDAP to import users or user groups from an LDAP server. LDAP users can use their LDAP credentials to log in to vRealize Operations Manager.
  - Create a single sign-on source and import users and user groups from a single sign-on server. Single sign-on users can use their single sign-on credentials to log in to vRealize Operations Manager and vCenter Server. You can also use Active Directory through single sign-on by configuring the Active Directory through single sign-on and adding the single sign-on source to vRealize Operations Manager.

User Preferences

To determine the display options for vRealize Operations Manager, such as colors for the display and health chart, the number of metrics and groups to display, and whether to synchronize system time with the host machine, you configure the user preferences on the top toolbar.

Users of vRealize Operations Manager

Each user has an account to authenticate them when they log in to vRealize Operations Manager.

The accounts of local users and LDAP users are visible in the vRealize Operations Manager user interface when they are set up. The accounts of vCenter Server and single sign-on users only appear in the user interface after a user logs in for the first time. Each user can be assigned one or more roles, and can be an authenticated member of one or more user groups.

Local Users in vRealize Operations Manager

When you create user accounts in a local vRealize Operations Manager instance, vRealize Operations Manager stores the credentials for those accounts in its global database, and authenticates the account user locally.

Each user account must have a unique identity, and can include any associated user preferences.

If you are logging in to vRealize Operations Manager as a local user, and on occasion receive an `invalid password` message, try the following workaround. In the Login page, change the Authentication Source to **All vCenter Servers**, change it back to **Local Users**, and log in again.

vCenter Server Users in vRealize Operations Manager

vRealize Operations Manager supports vCenter Server users. To log in to vRealize Operations Manager, vCenter Server users must be valid users in vCenter Server.

Roles and Associations

A vCenter Server user must have either the vCenter Server Admin role or one of the vRealize Operations Manager privileges, such as PowerUser which assigned at the root level in vCenter Server, to log in to vRealize Operations Manager. vRealize Operations Manager uses only the vCenter privileges, meaning the vRealize Operations Manager roles, at the root level, and applies them to all the objects to which the user has access. After logging in, vCenter Server users can view all the objects in vRealize Operations Manager that they can already view in vCenter Server.

Logging in to vCenter Server Instances and Accessing Objects

vCenter Server users can access either a single vCenter Server instance or multiple vCenter Server instances, depending on the authentication source they select when they log in to vRealize Operations Manager.

- If users select a single vCenter Server instance as the authentication source, they have permission to access the objects in that vCenter Server instance. After the user has logged in, an account is created in vRealize Operations Manager with the specific vCenter Server instance serving as the authentication source.

- If users select **All vCenter Servers** as the authentication source, and they have identical credentials for each vCenter Server in the environment, they see all the objects in all the vCenter Server instances. Only users that have been authenticated by all the vCenter Servers in the environment can log in. After a user has logged in, an account is created in vRealize Operations Manager with all vCenter Server instances serving as the authentication source.

vRealize Operations Manager does not support linked vCenter Server instances. Instead, you must configure the vCenter Server adapter for each vCenter Server instance, and register each vCenter Server instance to vRealize Operations Manager.

Only objects from a specific vCenter Server instance appear in vRealize Operations Manager. If a vCenter Server instance has other linked vCenter Server instances, the data does not appear.

vCenter Server Roles and Privileges

You cannot view or edit vCenter Server roles or privileges in vRealize Operations Manager. vRealize Operations Manager sends roles as privileges to vCenter Server as part of the vCenter Server Global privilege group. A vCenter Server administrator must assign vRealize Operations Manager roles to users in vCenter Server.

vRealize Operations Manager privileges in vCenter Server have the role appended to the name. For example, vRealize Operations Manager ContentAdmin Role, or vRealize Operations Manager PowerUser Role.

Read-Only Principal

A vCenter Server user is a read-only principal in vRealize Operations Manager, which means that you cannot change the role, group, or objects associated with the role in vRealize Operations Manager. Instead, you must change them in the vCenter Server instance. The role applied to the root folder applies to all the objects in vCenter Server to which a user has privileges. vRealize Operations Manager does not apply individual roles on objects. For example, if a user has the PowerUser role to access the vCenter Server root folder, but has read-only access to a virtual machine, vRealize Operations Manager applies the PowerUser role to the user to access the virtual machine.

Refreshing Permissions

When you change permissions for a vCenter Server user in vCenter Server, the user must log out and log back in to vRealize Operations Manager to refresh the permissions and view the updated results in vRealize Operations Manager. Alternatively, the user can wait for vRealize Operations Manager to refresh. The permissions refresh at fixed intervals, as defined in the `$ALIVE_BASE/user/conf/auth.properties` file. The default refreshing interval is half an hour. If necessary, you can change this interval for all nodes in the cluster.

Single Sign-On and vCenter Users

When vCenter Server users log into vRealize Operations Manager by way of single sign-on, they are registered on the vRealize Operations Manager User Accounts page. If you delete the account of a vCenter Server user that has logged into vRealize Operations Manager by way of single sign-on, or remove the user from a single sign-on group, the user account entry still appears on the User Account page and you must delete it manually.

Generating Reports

vCenter Server users cannot create or schedule reports in vRealize Operations Manager.

Backward Compatibility for vCenter Server Users in vRealize Operations Manager

vRealize Operations Manager provides backward compatibility for users of the earlier version of vRealize Operations Manager, so that users of vCenter Server who have privileges in the earlier version in vCenter Server can log in to vRealize Operations Manager.

When you register vRealize Operations Manager in vCenter Server, certain roles become available in vCenter Server.

- The Administrator account in the previous version of vRealize Operations Manager maps to the PowerUser role.
- The Operator account in the previous version of vRealize Operations Manager maps to the ReadOnly role.

During registration, all roles in vRealize Operations Manager, except for vRealize Operations Manager Administrator, Maintenance, and Migration, become available dynamically in vCenter Server. Administrators in vCenter Server have all of the roles in vRealize Operations Manager that map during registration, but these administrator accounts only receive a specific role on the root folder in vCenter Server if it is specially assigned.

Registration of vRealize Operations Manager with vCenter Server is optional. If users choose not to register vRealize Operations Manager with vCenter Server, a vCenter Server administrator can still use their user name and password to log in to vRealize Operations Manager, but these users cannot use the vCenter Server session ID to log in. In this case, typical vCenter Server users must have one or more vRealize Operations Manager roles to log in to vRealize Operations Manager.

When multiple instances of vCenter Server are added to vRealize Operations Manager, user credentials become valid for all of the vCenter Server instances. When a user logs in to vRealize Operations Manager, if the user selects all vCenter Server options during login, vRealize Operations Manager requires that the user's credentials are valid for all of the vCenter Server instances. If a user account is only valid for a single vCenter Server instance, that user can select the vCenter Server instance from the login drop-down menu to log in to vRealize Operations Manager.

vCenter Server users who log in to vRealize Operations Manager must have one or more of the following roles in vCenter Server:

- vRealize Operations Content Admin Role
- vRealize Operations General User Role 1
- vRealize Operations General User Role 2
- vRealize Operations General User Role 3
- vRealize Operations General User Role 4
- vRealize Operations Power User Role

- vRealize Operations Power User without Remediation Actions Role
- vRealize Operations Read Only Role

For more information about vCenter Server users, groups, and roles, see the vCenter Server documentation.

External User Sources in vRealize Operations Manager

You can obtain user accounts from external sources so that you can use them in your vRealize Operations Manager instance.

There are two types of external user identity sources:

- Lightweight Directory Access Protocol (LDAP): Use the LDAP source if you want to use the Active Directory or LDAP servers as authentication sources. The LDAP source does not support multi-domains even when there is a two-way trust between Domain A and Domain B.
- Single Sign-On (SSO): Use a single sign-on source to perform single sign-on with any application that supports vCenter single sign-on, including vRealize Operations Manager. For example, you can install a standalone vCenter Platform Services Controller (PSC) and use it to communicate with an Active Directory server. Use a PSC if the Active Directory has a setup that is too complex for the simple LDAP source in vRealize Operations Manager, or if the LDAP source is experiencing slow performance.

Roles and Privileges in vRealize Operations Manager

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

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

Each predefined role includes a set of privileges for users to perform create, read, update, or delete actions on components such as dashboards, reports, administration, capacity, policies, problems, symptoms, alerts, user account management, and adapters.

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.

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 new 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 vRealize Operations Manager Authentication Sources .

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 vRealize Operations Manager Cluster and Node Maintenance.
- Verify that you have access to a Platform Services Controller through the vCenter Server. See the VMware vSphere Information Center for more details.

Procedure

1. Log in to vRealize Operations Manager as an administrator.
2. In the menu, click **Administration**, then in the left pane click **Access > Authentication Sources**.
3. Click **Add**.
4. In the Add Source for User and Group Import dialog box, provide information for the single sign-on source.

Option	Action
Source Display Name	Type a name for the import source.
Source Type	Verify that SSO SAML is displayed.
Host	Enter the IP address or FQDN of the host machine where the single sign-on server resides. If you enter the FQDN of the host machine, verify that every non-remote collector node in the vRealize Operations Manager cluster can resolve the single sign-on host FQDN.
Port	Set the port to the single sign-on server listening port. By default, the port is set to 443.
User Name	Enter the user name that can log into the SSO server.
Password	Enter the password.
Grant administrator role to vRealize Operations Manager for future configuration?	Select **Yes** so that the SSO source is reregistered automatically if you make changes to the vRealize Operations Manager setup. If you select **No**, and the vRealize Operations Manager setup is changed, single sign-on users will not be able to log in until you manually reregister the single sign-on source.
Automatically redirect to vRealize Operations single sign-on URL?	Select **Yes** to direct users to the vCenter single-sign on log in page. If you select **No**, users are not redirected to SSO for authentication. This option can be changed in the vRealize Operations Manager Global Settings.

Option	Action
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 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 4-167. Access Control Tabs and Workspaces

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 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 he can perform those actions 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 and Workspaces

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 4-168. Access Control User Accounts Summary Grid

Summary Grid Options	Description
User Accounts toolbar	To manage user accounts, use the toolbar icons. ■ **Add** icon. Add a user account, and provide the details for the user account in the Add User Account workspace. ■ **Edit** icon. Edit the selected user account, and modify the details for the user group in the Edit User Account workspace. ■ **Delete** icon. Delete a user account. ■ **Import Users** icon. 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.
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.

Table 4-168. Access Control User Accounts Summary Grid (continued)

Summary Grid Options	Description
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 could become 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 of 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 4-169. 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. ■ Group Name: Identifies the user group. To change the user groups associated with the user account, click the **Edit** icon. ■ 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.

User Accounts Add or Edit User Workspace: User Details

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.

Where You Add or Edit User Accounts

To add a user account, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.

Table 4-170. Access Control Add or Edit User Workspace - User Details Page

User Details Options	Description
User Name	User name, without spaces, that will log in to 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 to require users to change their password the next time they log in to the vRealize Operations Manager instance.

Table 4-171. Access Control Add or Edit User Workspace - 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** checkbox. 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** When you assign a user permission to take action on a parent object, such as an adapter, that user can perform the same action on all the parent's child objects. For example, if a user has permission to access the vRealize Operations Manager adapter, that user can access all the virtual machines associated with the adapter. This is true even if the same user holds another role that permits limited access to only one specific virtual machine.

Add or Edit User Workspace for User Accounts: Assign Groups, Roles, and Objects

You can assign a user account 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 Assign Groups, Roles, and Objects to User Accounts

To assign groups, roles, and objects to a user account, click **Administration**, and then in the left pane click **Access > Access Control**.

Table 4-172. Access Control Add or Edit User Workspace - 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** checkbox. 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** When you assign a user permission to take action on a parent object, such as an adapter, that user can perform the same action on all the parent's child objects. For example, if a user has permission to access the vRealize Operations Manager adapter, that user can access all the virtual machines associated with the adapter. This is true even if the same user holds another role that permits limited access to only one specific virtual machine.

Import Users Workspace for User Accounts: 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.

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 **Import Users** icon on the User Accounts toolbar.

Table 4-173. Access Control Import Users Workspace - Import Users Page

User Details Options	Description
Import From	LDAP host machine configured as the source to import the 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. 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.

Import Users Workspace for User Accounts: Assign Groups, Roles, and Objects

When you import a user account to vRealize Operations Manager, you assign the user account to user groups, assign roles, and specify the objects the user account can access when assigned each role.

Where You Assign Groups, Roles, and Objects to Imported User Accounts

1 To assign groups, roles, and objects to an imported user accounts, click **Administration**, and then in the left pane click **Access > Access Control**.

2 Click the **Import Users** icon on the User Accounts toolbar.

Table 4-174. Access Control Import Users Workspace - 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 and Workspace

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**, then in the left pane click **Access > Access Control**.

2 Click the **User Groups** tab.

Table 4-175. Access Control User Groups Summary Grid

Option	Description
User Groups toolbar	To manage user groups, use the toolbar icons. ■ **Add** icon. Add a user group, and provide the details for the user group in the Add User Group workspace. ■ **Edit** icon. Edit the selected user group, and modify the details for the user group in the Edit User Group workspace. ■ **Clone Group** icon. Clone a user group, and type a name and description for the cloned user group. ■ **Delete** icon. Delete a user group. ■ **Import Group** icon. Import a user group, and provide the details to import the user group in the Import User Groups workspace.
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 of 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 4-176. Access Control User Groups Details Grid

Option	Description
User Accounts	You can add members to the selected group, view only the selected or deselected members in the group, or search for a member. You can remove a user from the group by selecting the user in the Details pane and clicking **Delete**. ■ User Name: Name of each user who is a member of the selected group. ■ First Name: First name of each user in the group. ■ Last Name: Last name of each user in the group.
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.

Access Control: User Groups Add or Edit User Group

You can view and modify the details for user groups, including users, roles and objects.

Where You Add or Edit User Groups

1 To add a user group, in the menu, click **Administration**, then click **Access > Access Control**.

2 Click the **Add**.

3 To edit a user group, select a user group and click the **Edit** icon.

Table 4-177. Add or Edit User Group - Name and Description

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.

Table 4-178. 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.

Access Control: Import User Groups

You import user groups from a single sign-on server, 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 add import a user group, in the menu, click **Administration**, then in the left pane click **Access > Access Control**.

2 Click the **Import Group**.

3 To edit a user group, select a user group and click the **Edit** icon.

The options displayed in the Import User Groups page depend upon the authentication source you select.

When you import a user group from a single sign-on server, log out of vRealize Operations Manager, and then log in again to synchronize users and user group memberships with the single-sign on server.

Table 4-179. Import User Groups Workspace - Import User Groups Page - LDAP Source Options

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 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.

Table 4-179. Import User Groups Workspace - Import User Groups Page - LDAP Source Options (continued)

Option	Description
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 type 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 type 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 4-180. Import User Groups Workspace - Import User Groups Page - Single Sign-On Source Options

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 4-181. Import User Groups Workspace - 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 4-182. Access Control Roles Summary Grid

Option	Description
Roles toolbar	To manage roles, use the toolbar icons. ■ **Add** icon. Add a user role, and provide the name and description for the role in the Create Role dialog box. ■ **Edit** icon. Edit the selected user role, and modify the details for the role in the Edit Role dialog box. ■ **Clone** icon. Clone the selected user role. ■ **Delete** icon. 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 4-183. 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 users 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.

The actions named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear, but 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 actions named `Set Memory for VM Power Off Allowed`, `Set CPU Count for VM Power Off Allowed`, and `Set CPU Count and Memory for VM Power Off Allowed` are also not visible except in the alert recommendations, and 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 4-184. 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 checked, 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 checked, users must include one or more uppercase characters. ■ Passwords must contain special characters. When checked, 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.

vRealize Operations Manager Authentication Sources

vRealize Operations Manager uses two 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, and single sign-on.

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 4-185. 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. ■ **Edit** icon: Edit the selected authentication source, and modify the details in the Edit Source dialog box. ■ **Delete** icon. Delete an authentication source. ■ **Synchronize User Groups** icon. Synchronize LDAP users in the selected LDAP user groups.
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. ■ 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 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.
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 4-186. 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 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 4-187. 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 4-187. 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 as soon as 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 4-188. 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. ■ 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 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. ■ 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. 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

Table 4-188. Authentication Sources Add Source for User and Group Import - options available when **Open LDAP**, **Active Directory**, and **Other** are selected. (continued)

Option	Description
	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. ■ 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.
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 type 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 type 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 4-189. 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 vIDM 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.
Username	vIDM system-domain tenant administrator username.
Password	Password of the vIDM system-domain tenant administrator.
Redirect IP	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 master node. **Note** When the master replica becomes the master 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 master node.
Test	Tests whether the vIDM 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 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 startup and shutdown, application failures, watchdog restarts, configuration changes of applications, changes to security policy, requests, responses, and status of success.
- You must validate that the components installed in your environment are running the latest version.

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 4-190. 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 checked, 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.
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.

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.

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 4-191. 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 4-192. 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 4-193. 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 4-194. User Preference Settings

Option	Description
Display	Configure how many metrics and root cause groups to display. ■ 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 secure machine-to-machine communication within vRealize Operations Manager itself or between vRealize Operations Manager and other systems.

Change the vRealize Operations Manager Administrator Password

You might need to change the vRealize Operations Manager administrator password as part of securing or maintaining your deployment.

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 username and password for the master node.
3. In the upper right, click the **admin** drop-down menu, and click **Change Administrator Password**.
4. Enter the current password, and enter the new password twice to ensure its accuracy.

   Note You cannot change the administrator username of admin.

5. Click **OK**.

Reset the vRealize Operations Manager Administrator Password on vApp Clusters

You need to 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 master 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
```

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 master administrator login credentials, which might be a security issue.

A temporary passphrase is good for one use only.

Prerequisites

Create and configure the master 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 username and password for the master node.
3. In the list of cluster nodes, select the master node.
4. From the toolbar above the list, click the option to generate a passphrase.
5. Enter a number of hours before the passphrase expires.

6 Click **Generate**.

A random alphanumeric string appears, which you can send to a user who needs to add a node.

What to do next

Have the user supply the passphrase when adding a node.

Custom vRealize Operations Manager Certificates

By default, vRealize Operations Manager includes its own authentication certificates. The default certificates cause the browser to display a warning when you connect to the vRealize Operations Manager user interface.

Your site security policies might require that you use another certificate, or you might want to avoid the warnings caused by the default certificates. In either case, vRealize Operations Manager supports the use of your own custom certificate. You can upload your custom certificate during initial master node configuration or later.

Custom vRealize Operations Manager Certificate Requirements

A certificate used with vRealize Operations Manager must conform to certain requirements. Using a custom certificate is optional and does not affect vRealize Operations Manager features.

Requirements for Custom Certificates

Custom vRealize Operations Manager certificates must meet the following requirements.

- The certificate file must include the terminal (leaf) server certificate, a private key, and all issuing certificates if the certificate is signed by a chain of other certificates.
- In the file, the leaf certificate must be first in the order of certificates. After the leaf certificate, the order does not matter.
- In the file, all certificates and the private key must be in PEM format. vRealize Operations Manager does not support certificates in PFX, PKCS12, PKCS7, or other formats.
- In the file, all certificates and the private key must be PEM-encoded. vRealize Operations Manager does not support DER-encoded certificates or private keys.

  PEM-encoding is base-64 ASCII and contains legible BEGIN and END markers, while DER is a binary format. Also, file extension might not match encoding. For example, a generic `.cer` extension might be used with PEM or DER. To verify encoding format, examine a certificate file using a text editor.
- The file extension must be `.pem`.
- The private key must be generated by the RSA or DSA algorithm.
- The private key must not be encrypted by a pass phrase if you use the master node configuration wizard or the administration interface to upload the certificate.
- The REST API in this vRealize Operations Manager release supports private keys that are encrypted by a pass phrase. Contact VMware Technical Support for details.

- The vRealize Operations Manager Web server on all nodes will have the same certificate file, so it must be valid for all nodes. One way to make the certificate valid for multiple addresses is with multiple Subject Alternative Name (SAN) entries.
- SHA1 certificates creates browser compatibility issues. Therefore, ensure that all certificates that are created and being uploaded to vRealize Operations Manager are signed using SHA2 or newer.
- The vRealize Operations Manager supports custom security certificates with key length up to 8192 bits. An error is displayed when you try to upload a security certificate generated with a stronger key length beyond 8192 bits.

For more information, see the following KB articles:

- vRealize Operations Manager 6.x fails to accept and apply Custom CA Certificate (2144949)

Verifying a Custom vRealize Operations Manager Certificate

When you upload a custom certificate file, the vRealize Operations Manager interface displays summary information for all certificates in the file.

For a valid custom certificate file, you should be able to match issuer to subject, issuer to subject, back to a self-signed certificate where the issuer and subject are the same.

In the following example, `OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32` is issued by `OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32`, which is issued by `OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84`, which is issued by itself.

```
Thumbprint: 80:C4:84:B9:11:5B:9F:70:9F:54:99:9E:71:46:69:D3:67:31:2B:9C
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32
Subject Alternate Name:
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:24.000Z
Valid To: 2020-05-06T16:25:24.000Z

Thumbprint: 72:FE:95:F2:90:7C:86:24:D9:4E:12:EC:FB:10:38:7A:DA:EC:00:3A
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-
e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:19.000Z
Valid To: 2020-05-06T16:25:19.000Z

Thumbprint: FA:AD:FD:91:AD:E4:F1:00:EC:4A:D4:73:81:DB:B2:D1:20:35:DB:F2
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-
e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-
e1806f0d3f84
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
```

```
Valid From: 2015-05-07T16:24:45.000Z
Valid To: 2020-05-06T16:24:45.000Z
```

Sample Contents of Custom vRealize Operations Manager Certificates

For troubleshooting purposes, you can open a custom certificate file in a text editor and inspect its contents.

PEM Format Certificate Files

A typical PEM format certificate file resembles the following sample.

```
-----BEGIN CERTIFICATE-----
MIIF1DCCBLygAwIBAgIKFYXYUwAAAAAGTANBgkqhkiG9w0BAQ0FADBhMRMwEQYK
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 4-195. 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 4-196. 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 master node, you can still add a certificate after vRealize Operations Manager is installed.

Prerequisites

- Create and configure the master node.

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 username and password.
3. At the upper right, click the yellow 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 username 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 the **x** to remove the certificate.
6. If the certificate is being used by the adapter, then the following message comes up:

   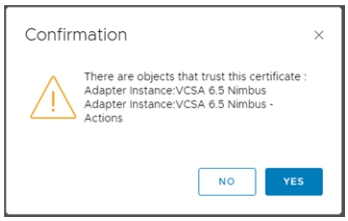

   A certificate can be configured for one or more adapters if it is the same destination system.
7. If you delete a certificate which is already being used by another adapter, the adapter fails to connect or start. As a workaround, perform the following steps:
   a. On the left pane, click **Solutions**.
   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.

Modifying Global Settings

The global settings control the system settings for vRealize Operations Manager, including data retention and system timeout settings. You can modify one or more of the settings to monitor your environment better. These settings affect all your users.

The global settings do not affect metric interactions, color indicators, or other object management behaviors. These behaviors are configured in your policies.

Settings related to managing objects with vRealize Operations Manager are available on the **Inventory Explorer** page.

You can view tooltips for each option in the Edit Global Settings dialog box.

Global Settings Best Practices

Most of the settings pertain to how long vRealize Operations Manager retains collected and process data.

The default values are common retention periods. You might need to adjust the time periods based on your local policies or disk space.

List of Global Settings

The global settings determine how vRealize Operations Manager retains data, keeps connection sessions open, and other settings. These are system settings that affect all users.

Table 4-197. 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 might be identified by vRealize Operations Manager as not existing and vRealize Operations Manager can no longer collect data about the object. Whether vRealize Operations Manager identifies deleted objects as not existing depends the adapter. This feature is not implemented in some adapters. For example, if the retention time is 360 hour and a virtual machine is deleted from a vCenter Server instance, the virtual machine remains as an object in vRealize Operations Manager for 15 days before it is deleted. This setting applies to objects deleted from the data source or server, not to any objects you delete from vRealize Operations Manager on the Inventory Explorer page. A value of **–1** deletes objects immediately.
Deletion Schedule Interval	24 hours	Determines the frequency to schedule deletion of resources. This setting works with the Deleted Objects setting to remove objects that no longer exist in the environment. vRealize Operations Manager transparently marks objects for removal that have not existed for the length of time specified under Deleted Objects. vRealize Operations Manager then removes the marked objects at the frequency specified under Deletion Scheduling Interval.
Object History	300 days	Number of days to retain the history of the object configuration, relationship, and property data. The configuration data is the collected data from the monitored objects on which the metrics are based. The collected data includes changes to the configuration of the object. The data is purged from the system after the specified number of days.
Session Timeout	30 minutes	If your connection to vRealize Operations Manager is idle for the specified amount of time, you are logged out of the application. You must provide credentials to log back in.
Symptoms/Alerts	45 days	Number of days to retain cancelled alerts and symptoms. The alerts and symptoms can be canceled by the system or canceled 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.

Table 4-197. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Additional Time Series Retention	36 months	The additional time series retention is always larger than the Time series data retention and the roll up data is available starting from the end of the regular period and till the end of the roll-up data retention period. If 0 value is specified then this will effectively disable Additional Time Series Retention and only data specified in Time Series Retention will be stored.
Dynamic Threshold Calculation	enabled	Determines whether to calculate normal levels of threshold violation for all objects. If the setting is disabled, the following areas of vRealize Operations Manager will not work or are not displayed: ■ Anomalies badge is not calculated ■ 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.
Capacity Calculation	enabled	Determines whether to calculate capacity metrics and badges for all objects. If the setting is disabled, the values for the following badges are not calculated: ■ Capacity Remaining ■ Time Remaining ■ Stress ■ Reclaimable Capacity ■ Density
Allow vCenter users to log in to individual vCenters using the vRealize Operations Manager UI		Determine how users of vCenter Server log in to vRealize Operations Manager. ■ In the vRealize Operations Manager user interface, vCenter Server users can log in to individual vCenter Server instances. Disabled by default. ■ vCenter Server users can log in from vCenter Server clients. Enabled by default. ■ In the vRealize Operations Manager user interface, vCenter Server users can log in to all vCenter Server instances. Enabled by default.
Allow vCenter users to log in from vCenter clients		
Allow vCenter users to log in to all vCenters using the vRealize Operations Manager UI		

Table 4-197. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Automated Actions	enabled or disabled	Determines whether to allow vRealize Operations Manager to automate actions. When an alert triggers, 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.
Enable Certificate Validation Standard Way		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.
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.

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 left pane click **Management > Global Settings**.

2 To edit the global settings, click the Edit icon.

Table 4-198. Global Settings Options

Option	Description
Edit Global Settings	Use the toolbar option to modify setting values.
Setting	Setting name.
Value	Current value for the setting. To change the setting value, click **Edit Global Settings**.
Description	Information about the setting. Place your mouse over 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.

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 .

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**.

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 and limit the retained log size, send the logs to an external syslog server, and set logging levels.

vRealize Operations Manager logs are categorized by cluster node, and log type.

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 are listed under the node.

2 Select the **Syslog Export** checkbox.

3 Click the gear icon, enter the **IP address or Host Name**, and **Port**.

4 Click **OK**.

Note Not all components have relevant syslog information. Therefore, not all nodes have the configuration option enabled.

Figure 4-1. Logs

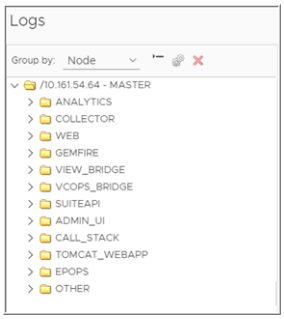

Figure 4-2. Log Options

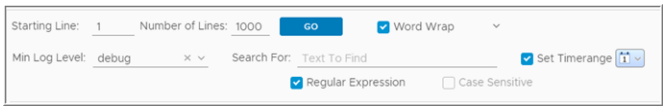

Table 4-199. Log Viewer Toolbar Options

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.

Table 4-199. Log Viewer Toolbar Options (continued)

Option	Description
Edit Properties	For the selected folder, you can limit the log size, send the logs to an external syslog server, and set logging levels. **Caution** The logs that you transmit to a syslog server are unencrypted. Verify that your network is secure before using the syslog option.
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.
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 cursor 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.

Use the toolbar options to add, download, or remove items.

Table 4-200. 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.

Table 4-200. Support Bundle Toolbar Options (continued)

Option	Description
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 4-201. 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 4-202. 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

Table 4-202. Dynamic Threshold Options (continued)

Option	Description
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, as well as 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 4-203. Adapter Redescribe Options

Option	Description
Redescribe	Start the adapter describe process

vRealize Operations Manager provides adapter-specific details from the redescribe process.

Table 4-204. 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 Managerobtains 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 4-205. 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 **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 tovRealize 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 4-206. 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

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.

- You need to enable or disable high availability (HA).
- You need to upload and install vRealize Operations Manager software update PAK files.
- The product interface is inaccessible, and you need to correct the problem by bringing nodes online, or by restarting nodes or the cluster.
- vRealize Operations Manager needs to be restarted for any reason.

Note that 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.

vRealize Operations Manager Cluster Status and Troubleshooting

vRealize Operations Manager includes a central page where you can monitor and manage the nodes in your vRealize Operations Manager cluster as well as the adapters that are installed on the nodes.

How Cluster Status and Troubleshooting Works

Cluster status and troubleshooting lets you 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 Status and Troubleshooting

Log in to the vRealize Operations Manager administration interface at https://*master-node-name-or-ip-address*/admin.

Cluster Status and Troubleshooting Options

The options include cluster-level monitoring and management features.

Table 4-207. 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.
High Availability	Indicates whether HA 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 4-208. 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. Master and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.
Cluster Role	Type of vRealize Operations Manager node: master, 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.
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

In addition, there are adapter statistics for the selected node.

Table 4-209. 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 4-210. 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 4-211. 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 4-212. 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

Configuring and Using Workload Balance

Workload Balance is a feature that provides for moving virtual compute resources and their associated file systems dynamically across datastore clusters within a data center or custom data center.

Using Workload Balance, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining and/or improving cluster and datacenter performance.

Workload Balance further enables you to potentially automate a significant portion of your datacenter compute and storage balancing efforts. With properly defined policies determining the threshold at which resource contention triggers an alert and automatically executes a recommended action, a datacenter can perform at optimum.

Configuring Workload Balance

Workload Balance offers you the potential to fully automate a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

1 Set the policies that define demand limits for cluster CPU and memory. See Workload Policy Settings.

2 Configure the Workload Automation Details to complement the policies. See Workload Automation Details.

3 Configure the two Workload Balance alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the recommended actions calculated by Workload Balance are executed automatically. See Configuring Workload Balance Alerts

Prerequisites

Workload Balance 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 recommends that you rebalance clusters in a data center, the system does not guarantee a rebalance action will occur. vRealize Operations Manager analytics can determine that a rebalance is desirable and can create a rebalancing plan. However, the system cannot automatically identify all of the architectural constraints that may be present. Such constraints may prevent a rebalance 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 via shared storage.
- Virtual machines defined with affinity rules and/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. This means, for example, that if a virtual machine that has a vmdk on a datastore and a second vmdk 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 vmdks on multiple datastores inside a single datastore cluster.
- Virtual machines cannot be placed inside clusters managed by vRealize Automation.
- Workload Balance may recommend moving virtual machines that are protected via vSphere Replication or Array Based Replication. It is recommended to ensure that all the clusters within a selected Datacenter or Custom Datacenter have replication available. You can set up DRS Affinity Rules on virtual machines that you do not want moving across clusters.

Workload Policy Settings

The key element in effecting optimum workload balance is appropriately setting the policies that determine when alerts are fired to signal that performance issues are detected. And the issues can be alleviated through a rebalance of system resources. For information on setting policies in vRealize Operations Manager see the following:

- Policies
- Default Policy in vRealize Operations Manager
- Policy Library Tab for Policies
- Policy Alert Definitions
- Using the Monitoring Policy Workspace to Create and Modify Operational Policies

Key among a number of settings relevant to the cluster workload are acceptable demand for CPU and memory. Diskspace demand is not considered as part of the workload balance equation.

In setting policies at the Add/Edit Monitoring Policy workspace, you arrive at a workload score that determines when demands on computing and memory are within an acceptable zone (green), moving toward contention (orange) or unacceptably high (red).

Note If a cluster CPU or memory utilization has reached the maximum configured capacity, there is no room left to move compute resources, and rebalancing is not possible. Set workload scores to values that trigger an alert well before absolute maximum capacity is reached.

Cluster Headroom Policy and Storage vMotion

The Add/Edit Monitoring Policy workspace includes additional options that more generally define how and when rebalance actions occur: click Workload Automation Details to set Balance, Consolidate, Headroom and Advanced Settings.

The Headroom option can potentially act more directly on rebalancing actions than the Balance or Consolidate workload options, because it sets a value that directly impacts the cluster workload limit policy. Headroom sets a buffer, that is, a workload limit separate from the workload score you defined as within the green zone in your policy settings. If your headroom buffer is, for example, 20 percent, the workload limit is effectively 80 percent, even if you set a workload score in the policy settings that permits the green zone to reach 90 percent. When Workload Balance calculates a rebalancing plan, it abides by the more restrictive metric.

Note Under Advanced Settings, you must ensure Storage vMotion is set to ON - the default.

Configuring Workload Balance Alerts

vRealize Operations Manager provides two preconfigured alerts designed to work with the Workload Balance feature. You must take additional action in the Policies area to turn on the alerts and automate them so that recommended actions are executed when the alerts fire.

The following preconfigured alerts are designed to work with the Workload Balance feature:

- Data center may require rebalance to relieve heavy resource utilization in one or more clusters
- Custom data center may require rebalance to relieve heavy resource utilization in one or more clusters

Prerequisites

Ensure you have all required permissions to access the Workload Balance UI 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 datacenters and custom datacenters, for example, **vSphere Solution's Default Policy**.
3. Click **Edit**.
4. Click the object grouping containing the relevant datacenters, for example, **vCenter Adapter Cluster Compute Resource**:

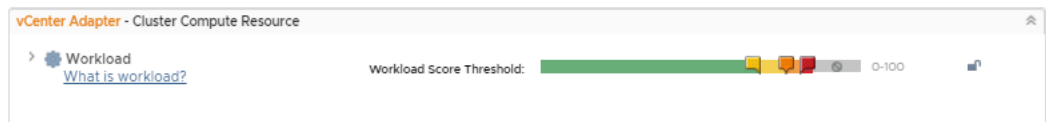

5. Click #6 on the lower left, Alert/Symptom Definitions.
6. Search on "may require rebalance" to locate the two alerts you want.
7. The alerts are turned ON by default/inheritance (State column).

8 The alerts are not automated by default/inheritance (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 Balance should be fully automated for your environment.

What to do next

Monitor rebalance activity at the User Scenario: Run Workload Rebalance to confirm that actions are taken automatically.

Using Workload Balance

Use the Workload Balance UI to monitor rebalancing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and execute 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 Balance screen, which is listed in the left pane on the Home screen. Depending on what appears on the screen, you may determine that you need to use rebalancing functions to distribute workload more evenly in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to identify if any active 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 Chapter 5 Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager.

The following user scenarios provide examples of the primary ways you can use Workload Balance to keep your data centers balanced and performing at optimum.

User Scenario: Run Workload Rebalance

As a virtual infrastructure administrator or other IT professional, you use Workload Balance functions to identify points of resource contention. In this scenario, you manually run a rebalance action to relieve CPU demand.

When you log into vRealize Operations Manager Home, you see the Recommended Actions page. In the Select Object Type pane, the Data Center column appears highlighted. Under the Worst Health heading, CMBU_ESO_VC09_DC appears, showing an associated alert. In the Suggested Fix table, the alert reads “Data Center may require rebalance to” and the suggested fix is “Rebalance the container to spread the workload more evenly.”

Prerequisites

Ensure that you have all required permissions to access the Workload Balance UI and manage vCenter Server objects.

Procedure

1 Click **Workload Balance** in the left pane.

Alternatively, you can hover the mouse over **Rebalance the container to spread the workload more evenly**. Select **Review the workload balance for the container**.

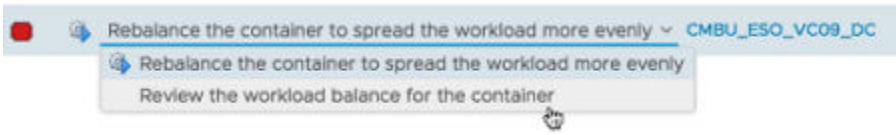

2 The Workload Balance page appears with the troubled Data Center preselected. If no data center is preselected, in the Select a Data Center to Review Workload pane, select a data center or custom data center from the list.

The panes and widgets to the right refresh to display data about the selected data center, CMBU_ESO_VC09_DC . Review the data points and trends shown.

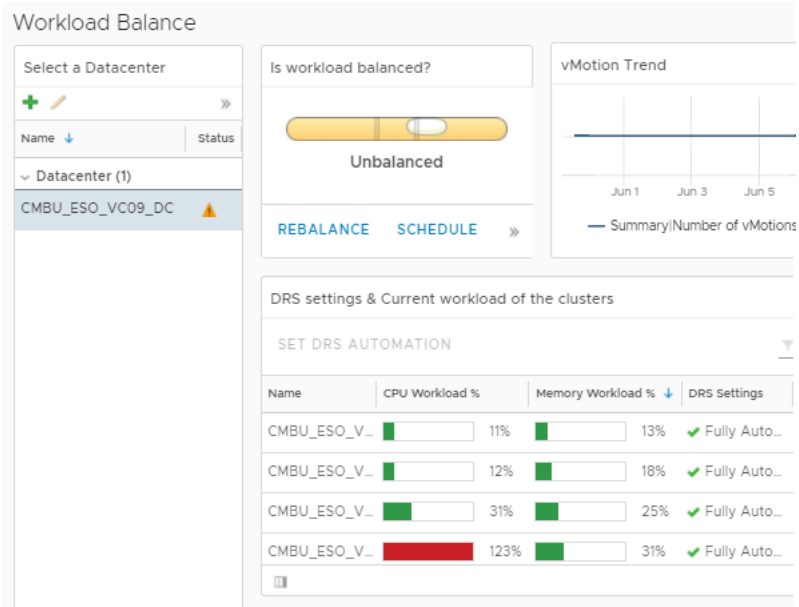

3 Based on the available data, you determine a rebalance action is required.

Table 4-213. Panes and Widgets

Pane	Contents
Is Workload Balanced?	A graphic of a carpenter's level appears in yellow, with the telltale air bubble off center, indicating an out-of-balance situation. (The air bubble is static, and does not convey varied out-of-balance states.)
What is the vMotion Trend?	No vMotions have happened in the last 24 hours.
DRS Settings and Current Workload	Shows where the imbalance arises: CPU workload on Cluster _01 is running at 107 percent of capacity. For a complete description of the widget, see the DRS Cluster Settings Widget.

Table 4-213. Panes and Widgets (continued)

Pane	Contents
Capacity Utilization of Datacenter Resources	In its Cluster Compute Resource graph, confirms that Cluster_01 CPU resources are running at 107 percent. For a complete description of the widget, see the Capacity Utilization Widget.
What is the Workload Trend?	Overall CPU usage for the data center has ticked upward slightly to 46 percent, with memory usage steady at 21 percent.

4 Click **REBALANCE** in the Is Workload Balanced? pane.

The system creates a rebalancing plan, which depicts BEFORE and (projected) AFTER workload statistics for the rebalance action.

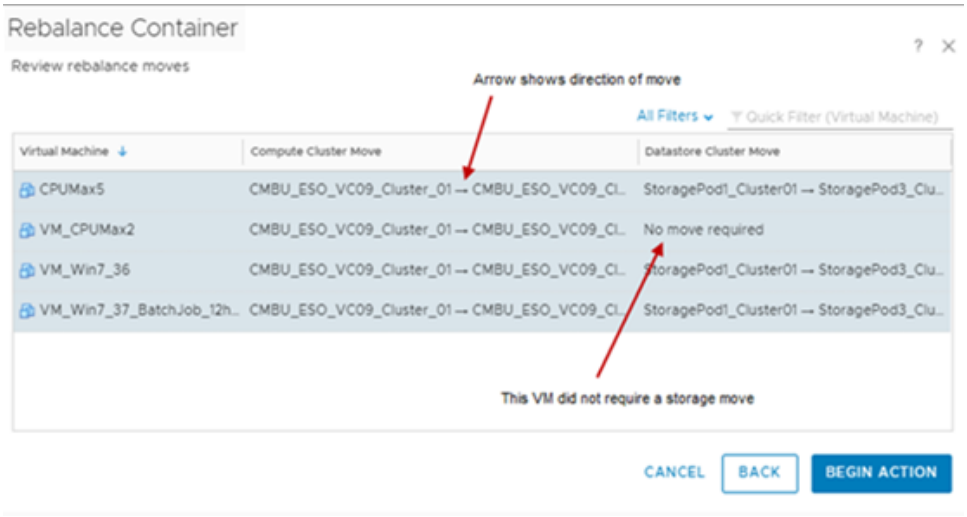

5 If you are satisfied with the projected results of the rebalance action, click **NEXT**. The dialog box updates to show the planned moves:

6 Review the rebalance moves, then click **BEGIN ACTION**.

The system runs the compute and storage resource moves.

Results

The rebalancing action moved compute and storage resources from an overloaded cluster to other clusters in the data center, thus reducing demand on any single cluster.

Note The Workload Balance page refreshes every five minutes. Depending on when you run a rebalancing 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 rebalance action was finished, go to the Recent Tasks page by selecting **Administration** on the menu, and clicking **History > Recent Task** 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 and enter the name of one of the VMs in the rebalance plan.

User Scenario: Schedule a Repeating Rebalancing Action

As a virtual infrastructure administrator or other IT professional, you determine that compute and storage resources in a given data center are unpredictable and a regularly scheduled rebalancing action would 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 Balance screen. Depending on what appears, you may determine that you need to schedule rebalancing functions to distribute workload more evenly in a data center or custom data center.

Prerequisites

Ensure you have all required permissions to access the Workload Balance UI and manage vCenter Server objects.

Procedure

1. From the Home screen, click **Workload Balance** in the left pane.
2. In the Select a Data Center to Review Workload pane, select a data center for which you want to schedule repeated rebalance actions.

3 In the Is Workload Balanced? pane, click **SCHEDULE**.

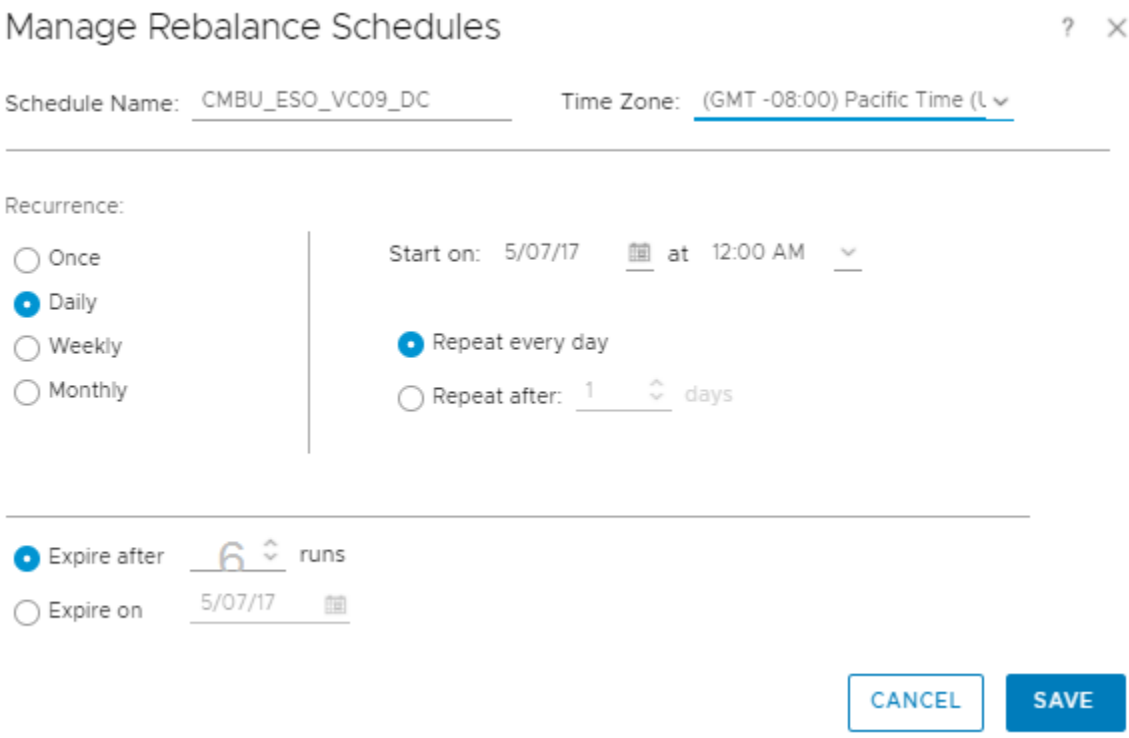

4 Give the schedule a name and choose a time zone.

5 Determine how often you want to repeat the rebalance 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 rebalance 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 rebalance action repeats for six days, then stops.

At the Workload Balance page, the Scheduled button appears in the upper right of the Are clusters balanced? pane as long as rebalance actions are scheduled for the selected data center. If you want to edit or delete a schedule, click the **Scheduled** button. The Using Rebalance Schedules page appears, where you can perform those actions.

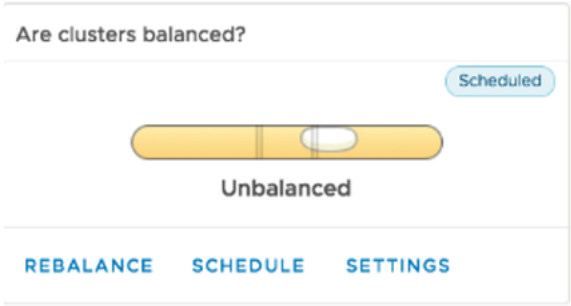

Note If you schedule a number of rebalancing actions close together, and the rebalance plans of two or more actions include overlapping functions, that is, they impact the same set of resources, the system shifts the actions into a queue. As a result, some actions may complete later than expected, with longer running actions and other potential system constraints extending the lag time. Rebalancing actions that do not overlap can run concurrently.

What to do next

To confirm that your rebalance action was finished, go to the Recent Tasks screen by selecting **Administration** on the 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 rebalance plan.

Note Because real-time data center resource contention is dynamic, the system calculates a new rebalancing plan each time the scheduled rebalancing action starts, but before it executes. If the system determines that the data center container is balanced at this moment, then the action is not executed. On the Recent Tasks page, the name of the affected data center appears in the Object Name column, and the Message “The balance of the selected container cannot be improved” appears under Details . Another possibility is that a scheduled rebalancing 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.

User Scenario: Execute Rebalance from the Home Screen

vRealize Operations Manager opens at Recommended Actions, with data center and custom data center errors highlighted. If a recommended rebalance action is available, it appears in the bottom third of the screen, with details.

To execute the action, click the blue **Run Action** arrow.

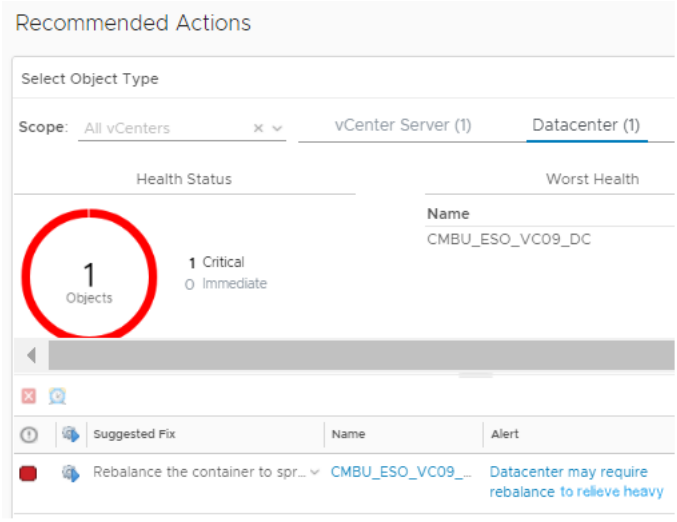

Prerequisites

Ensure you have all required permissions to access the Workload Balance UI and manage vCenter Server objects.

Results

The system executes the proposed rebalancing action.

What to do next

The Workload Balance 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 to search on. If the action succeeded, the Event Source column shows Alert: <*alert name*>.

Workload Balance Page

Workload Balance enables you to rebalance virtual machines and storage across datastore clusters to reduce resource contention and maintain optimum system performance.

Where You Find Workload Balance

From the Home screen, select Workload Balance in the left pane.

Table 4-214. Workload Balance Options

Option	Description
Select a Datacenter pane	Select a datacenter or custom datacenter from the list and all data to the right refreshes with information for the selected object. You can add a new customer datacenter from here and edit it.
Is Workload Balanced?	Graphic of carpenter's level shows a basic assessment: ■ Green with bubble centered = balanced workload ■ Yellow with bubble off-center = unbalanced workload Three major Workload Balance functions accessed here: ■ REBALANCE - dialogs to executing rebalance actions ■ SCHEDULE - dialog enabling scheduling one or more rebalance actions ■ SETTINGS - shortcut to Policy area's Workload Automation Details configuration
What is the Workload Trend?	Shows current and historical CPU and memory usage for the selected datacenter, with a graph to depict the trend
DRS Settings and Current Workload	Data derived from the DRS Cluster Settings Widget shows CPU, memory workload, DRS setting for each cluster; 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. Provides the option to set DRS automation level for individual objects.
Capacity Utilization of Datacenter Resources	Data derived from the Capacity Utilization Widget. When you point to an object, vRealize Operations Manager displays a popup summary that displays the object name, the capacity used by the object, and the reason that the capacity resource is constrained. To display the analysis details about the capacity on the object so that you can further troubleshoot the problem, click Details. By default the objects are constrained by the most constrained metric.

See also User Scenario: Run Workload Rebalance

Rebalance Containers Dialog Box

A two-page dialog that provides information about rebalancing the workload of a selected container.

First page: The current workload ("before," for example, CPU 105%) and projected results ("after," for example storage utilization 45%) for a possible rebalancing action.

Second page: The exact moves planned for compute and storage resources.

Where You Find Rebalance Clusters

At the Workload Balance screen, select REBALANCE from the Is Workload Balanced? pane.

Table 4-215. Rebalance Clusters Options

Option	Description
Compare Cluster Balance	If you are satisfied with the before and after numbers (First page, above), click NEXT.
Review Rebalance Moves	If you are satisfied with the moves planned (Second page, above), click BEGIN ACTION.

See also User Scenario: Run Workload Rebalance.

Manage Rebalance Schedules Dialog Box

Enables you to set up regular schedule for rebalancing a selected container.

Where You Find Manage Rebalance Schedules

At the Workload Balance screen, select SCHEDULES from the Is Workload Balanced? pane.

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 rebalance 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 5th of the month.
Start on:	Day to start rebalance schedule.
Start at:	Time to start rebalance schedule.
Expire after:	Designate a set number of scheduled runs.
Expire on:	Designate an exact date for the actions to end.

See also User Scenario: Schedule a Repeating Rebalancing Action

Workload Automation Policy Settings

Provides options for refining policy settings specifically for Workload Balance.

Where You Find Workload Automation Settings

At the Workload Balance screen, select **SETTINGS** from the Is Workload Balanced? pane.

Alternatively, you can 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 .

Using Rebalance Schedules

Use the Rebalance Schedules page to edit or delete rebalance schedules that you set up in the Manage Rebalance Schedule Dialog Box at the Workload Balance main screen.

Where You Find Rebalance Schedules

- From the Home screen, select **Administration > Configuration > Rebalance Schedules**.
- At the Workload Balance Page page, select in the Select a Data Center to Review Workload pane the data center whose rebalance schedule you want to edit or delete. Then click **Scheduled** in the Are Clusters Balanced? pane.

Table 4-216. Rebalance Schedules Options

Option	Description
Edit icon	Select a schedule from the list, then click the **Edit** icon. The Manage Rebalance Schedules Dialog Box 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 User Scenario: Run Workload Rebalance

Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager

5

You can use vRealize Operations Manager to resolve problems that your customers raise, respond to alerts that identify problems before your customers report problems, and generally monitor your environment for problems.

When your customers experience performance problems and call you to resolve the problem, the data that vRealize Operations Manager collects and analyzes is presented to you in graphical forms so that you can compare and contrast objects, understand the relationship between objects, and determine the root cause of problems.

To manage your environment as a proactive rather than reactive administrator, you monitor and respond to alerts. A generated alert notifies you when objects in your environment are experiencing problems. If you resolve the problem based on the alert before your customers notice, then you avoid service interruptions.

You can investigate the problems that generate alerts or that result in calls by using the **Alerts**, **Events**, **Details**, and **Environment** tabs.

If you find the root cause of the problem, you might be able to resolve the problem by running an action. The actions make changes to objects in the target system, for example, the VMware vCenter Server system, from vRealize Operations Manager.

This chapter includes the following topics:

- What to Do When...
- Monitoring and Responding to Alerts
- Monitoring and Responding to Problems
- Running Actions from vRealize Operations Manager
- Viewing Your Inventory

What to Do When...

As a virtual infrastructure administrator, network operations center engineer, or other IT professional, you use vRealize Operations Manager to monitor objects in your environment so that you can ensure service to your customers and resolve any problems that occur.

Your vRealize Operations Manager administrator has configured vRealize Operations Manager to manage two vCenter Server instances that manage multiple hosts and virtual machines. It is your first day using vRealize Operations Manager to manage your environment.

- User Scenario: A User Calls With a Problem

  The vice president of sales telephones the help desk reporting that her virtual machine, VPSALES4632, is running slow. She is working on sales reports for an upcoming meeting and is running behind schedule because of the slow performance of her virtual machine.

- User Scenario: An Alert Arrives in Your Inbox

  You return from lunch to find an alert notification in your inbox. You can use vRealize Operations Manager to investigate and resolve the alert.

- User Scenario: You See Problems as You Monitor the State of Your Objects

  As you investigate your objects in the context of this scenario, vRealize Operations Manager provides details to help you resolve the problems. You analyze the state of your environment, examine current problems, investigate solutions, and take action to resolve the problems.

User Scenario: A User Calls With a Problem

The vice president of sales telephones the help desk reporting that her virtual machine, VPSALES4632, is running slow. She is working on sales reports for an upcoming meeting and is running behind schedule because of the slow performance of her virtual machine.

As a network operations engineer, you were just reviewing the morning alerts and did not see any problems with her virtual machine, so you begin troubleshooting the problem.

Procedure

1. Search for a Specific Object

   As a network operations engineer, you must locate the customer's virtual machine in vRealize Operations Manager so that you can begin troubleshooting the reported problem.

2. Review Alerts Related to Reported Problems

   The sales vice president reports degraded performance in a virtual machine. To determine if the virtual machine has any alerts indicating the cause, review alerts for the virtual machine.

3. Use the Troubleshooting Tabs to Investigate a Reported Problem

   To troubleshoot problems with the VPSALES4632 virtual machine, as an example, you evaluate the symptoms, examine time line information, consider events, and create metric charts to find the root cause of the problem.

Search for a Specific Object

As a network operations engineer, you must locate the customer's virtual machine in vRealize Operations Manager so that you can begin troubleshooting the reported problem.

You use vRealize Operations Manager to monitor three vCenter Server instances with a total of 360 hosts and 18,000 virtual machines. The easiest way to locate a particular virtual machine is to search for it.

Procedure

1. In the **Search** text box, located on the vRealize Operations Manager title bar, type the name of the virtual machine.

   The **Search** text box displays all the objects that contain the string you type in the text box. If your customer knows that her virtual machine name contains SALES, you can type the string and the virtual machine is included in the list.
2. Select the object in the list.

Results

The main pane displays the object name and the **Summary** tab. The left pane displays and the related objects, including the host system and vCenter Server instance.

What to do next

Look for alerts related to the reported problem for the object. See Review Alerts Related to Reported Problems.

Review Alerts Related to Reported Problems

The sales vice president reports degraded performance in a virtual machine. To determine if the virtual machine has any alerts indicating the cause, review alerts for the virtual machine.

Alerts on an object can give you an insight into problems beyond the specific problem reported by the user.

Prerequisites

Locate the customer's virtual machine so that you can review related alerts. See Search for a Specific Object.

Procedure

1. Click the **Summary** tab for the object generating alerts.

   The **Summary** tab displays active alerts for the object.
2. Review the top alerts for Health, Risk, and Efficiency.

   Top alerts identify the primary contributors to the current state of the object. Do any of them appear to contribute to the slow response time? For example, any ballooning or swapping alerts indicate that you must add memory to the virtual machine. Are any alerts related to memory contention? Contention can be an indicator that you must add memory to the host.

3 If the **Summary** tab does not include top problems that appear to explain the reported problem, click the **Alerts** tab.

   The Alerts tab displays all active alerts for the current object.

4 Review the alerts for problems that are similar to or contribute to the reported problem.

   a To view the active and canceled alerts, click **Status: Active** to clear the filter and display active and inactive alerts.

      The canceled alerts might provide information about the problem.

   b So that you can locate alerts generated on or before the time when your customer reported the problem, click the **Created On** column to sort the alerts .

   c To view alerts for the parent objects in the same list with the alert for the virtual machine, click **View From**, then select, for example, **Host System** under Parents.

      The system adds these object types to the list so that you can determine if alerts among the parent objects are contributing to the reported problem.

5 If you locate an alert that appears to explain the reported problem, click the alert name in the alerts list.

6 On the **Alert > Symptoms** tabs, review the triggered symptoms and recommendations to determine if the alert indicates the root cause of the reported problem.

What to do next

- If the alert appears to indicate the source of the problem, follow the recommendations and verify the resolution with your customer. For an example, see Run a Recommendation on a Datastore to Resolve an Alert.
- If you cannot locate the cause of the reported problem among the alerts, begin more in-depth troubleshooting. See Use the Troubleshooting Tabs to Investigate a Reported Problem.

Use the Troubleshooting Tabs to Investigate a Reported Problem

To troubleshoot problems with the VPSALES4632 virtual machine, as an example, you evaluate the symptoms, examine time line information, consider events, and create metric charts to find the root cause of the problem.

If a review of the alerts did not help you identify the cause of the problem reported for the virtual machine, use the Troubleshooting tabs: **Alert > Symptoms**, **Event > Timeline**, and **All Metrics** to troubleshoot the history and current state of the virtual machine.

Prerequisites

- Locate the object for which the problem was reported. See Search for a Specific Object.
- Review the alerts for the virtual machine to determine if the problem is already identified and recommendations made. See Review Alerts Related to Reported Problems.

Procedure

1 In the menu, click **Environment**, then click **Inventory** and select VPSALES4632 from the tree.

The main pane updates to display the object **Summary** tab.

2 Click the **Alerts** tab, click the **Symptoms** tab, and review the symptoms to determine if one of the symptoms is related to the reported problem.

Depending on how your alerts are configured, some symptoms might be triggered but not sufficient to generate an alert.

a Review symptom names to determine if one or more symptoms are related to the reported problem.

The Information column provides the triggering condition, trend, and current value. What are the most common symptoms that affect response time? Do you see any symptoms related to CPU or memory usage?

b Sort by the **Created On** date so that you can focus on the time frame in which your customer reported that the problem.

c Click the **Status: Active** filter button to disable the filter so that you can review active and inactive symptoms.

Based on symptoms, you think the problem is related to CPU or memory use. But you do not know if the problem is with the virtual machine or with the host.

3 Click the **Events > Timeline** tabs and review the alerts, symptoms, and change events over time that might help you identify common trends that are contributing to the reported problem.

a To determine if other virtual machines had symptoms triggered and alerts generated at the same time as your reported problem, click **View From > Peer**.

Other virtual machine alerts are added to the time line. If you see that multiple virtual machines triggered symptoms in the same time frame, then you can investigate parent objects.

b Click **View From** and select **Host System** from the Parent list.

The alerts and symptoms that are associated with the host on which the virtual machine is deployed are added to the time line. Use the information to determine if a correlation exists between the reported problem and the alerts on the host.

4 Click the **Events > Events** tab to view changes in the collected metrics for the problematic virtual machine that could direct you toward the cause of the reported problem.

a Manipulate the **Date Controls** option view event for the approximate time when your customer reported the problem.

b Use the Filters to filter on event criticality and status. Select the Symptoms options if you want to include these in your analysis.

c Click an **Event** to view the details about the event.

d Click **View From**, select **Host System** under Parents, and repeat the analysis.

Comparing events on the virtual machine and the host, and evaluating those results, indicates that CPU or memory issues are the likely cause of the problem.

5 If you can identify that the problem is related to, for example, CPU or memory use, click the **All Metrics** tab to create your own metric charts so that you can determine whether it is one or the other, or a combination.

a If host is still the focus, then start by working with host metrics.

b In the metric list, double-click the **CPU Usage (%)** and the **Memory Usage (%)** metrics to add them to the workspace on the right.

c In the map, click the **VPSALES4632** object.

The metric list now displays the virtual machine metrics.

d In the metric list, double-click the **CPU Usage (%)** and the **Memory Usage (%)** metrics to add them to the workspace on the right.

e Review the host and virtual machine charts to see if you can identify a pattern that indicates the cause of the reported problem.

In this scenario, comparing the four charts reveals that CPU use is normal on both the host and the virtual machine, and the memory use is normal on the virtual machine. However, the memory use on the host began going consistently high three days before the reported problem on the VPSALES4632 virtual machine.

Results

The host memory is running consistently high, affecting the response time for the virtual machines. The number of virtual machines it is running is well within the supported amounts. The possible cause might be too many high process applications on the virtual machines. You can move some of the virtual machines to other hosts, distribute the workload, or power off idle virtual machines.

What to do next

- In this example, you can use vRealize Operations Manager to power off virtual machines on the host so that you can improve the performance of the virtual machines that are in use. See Run Actions From Toolbars in vRealize Operations Manager.
- If the combination of charts that you created on the **All Metrics** tab are something that you might want to use again, click **Generate Dashboard**.

User Scenario: An Alert Arrives in Your Inbox

You return from lunch to find an alert notification in your inbox. You can use vRealize Operations Manager to investigate and resolve the alert.

As a network operations engineer, you are responsible for several hosts and their datastores and virtual machines, and you receive emails when an alert is generated for your monitored objects. In addition to alerting you to problems in your environment, alerts should provide viable recommendations to resolve those problems. As you investigate this alert, you are evaluating the data to determine if one or more of the recommendations can resolve the problem.

This scenario assumes that you configured the outbound alerts to send standard email using SMTP and that you configured notifications to send you alert notifications using the standard email plug-in. When outbound alerts and notifications are configured, vRealize Operations Manager sends you messages when an alert is generated so that you can begin responding to problems as quickly as possible.

Prerequisites

- Verify that outbound alerts are configured for standard email alerts. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.
- Verify that the notifications are configured to send messages to your users for the alert definition. For an example of how to create an alert notification, see User Scenario: Create a vRealize Operations Manager Email Alert Notification .

Procedure

1 Respond to an Alert in Your Email

As a network operations engineer, you receive an email message from vRealize Operations Manager with information about one of the data stores for which you are responsible. The email notification informs you about the problem even when you are not presently working in vRealize Operations Manager.

2 Evaluate Other Triggered Symptoms for the Affected Data Store

Because you need more information about the data store before you decide on the best response, you examine the **Symptoms** tab to see other triggered symptoms for the data store.

3 Compare Alerts and Events Over Time in Response to a Datastore Alert

To evaluate an alert over time, compare the current alert and symptoms to other alerts and symptoms, other events, other objects, and over time.

4 View the Affected Datastore in Relation to Other Objects

To view the object for which the alert was generated as it relates to other objects, use the topological map on the **Relationships** tab.

5 Construct Metric Charts to Investigate the Cause of the Data Store Alert

To analyze the capacity metrics related to the generated alert, you create charts that compare different metrics. These comparisons help identify when something changed in your environment and what effect it had on the datastore.

6 Run a Recommendation on a Datastore to Resolve an Alert

As a network operations engineer, you investigated the alert regarding datastore disk space and determined that the provided recommendations can the problem. The recommendation to delete unused snapshots is especially useful. Use vRealize Operations Manager to delete the snapshots.

Respond to an Alert in Your Email

As a network operations engineer, you receive an email message from vRealize Operations Manager with information about one of the data stores for which you are responsible. The email notification informs you about the problem even when you are not presently working in vRealize Operations Manager.

In your email client, you receive an alert similar to the following message.

```
Alert was updated at Tue Jul 01 16:34:04 MDT :
Info:datastore1 Datastore is acting abnormally since Mon Jun 30 10:21:07 MDT  and was last updated at
Tue Jul 01 16:34:04 MDT

Alert Definition Name: Datastore is running out of disk space
Alert Definition Description: Datastore is running out of disk space
Object Name : datastore1
Object Type : Datastore
Alert Impact: risk
Alert State : critical
Alert Type : Storage
Alert Sub-Type : Capacity
Object Health State: info
Object Risk State: critical
Object Efficiency State: info
Symptoms:
SYMPTOM SET - self
Symptom Name    |  Object Name    |  Object ID    |  Metric    |  Message Info
Datastore space usage reaching critical limit    datastore1    |  b0885859-
e0c5-4126-8eba-6a21c895fe1b    |  Capacity|Used Space    |  HT above 99.20800922575977 > 95

Recommendations:
- Storage VMotion some Virtual Machines to a different Datastore
- Delete unused snapshots of Virtual Machines
- Add more capacity to the Datastore
Notification Rule Name: All alerts -- datastores
Notification Rule Description:
Alert ID : a9d6cf35-a332-4028-90f0-d1876459032b
Operations Manager Server - 192.0.2.0
Alert details
```

Prerequisites

- Verify that outbound alerts are configured for standard email alerts. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.

- Verify that the notifications are configured to send messages to your users for the alert definition. For an example of how to create an alert notification, see User Scenario: Create a vRealize Operations Manager Email Alert Notification .

Procedure

1 In your email client, review the message so that you understand the state of the affected objects and determine if you must begin investigating immediately.

  Look for the alert name, the alert state to determine the current level of criticality, and the affected objects.

2 In the email message, click **Alert Details**.

  vRealize Operations Manager opens on the **Summary** tab in the alert details for the generated alert and affected object.

3 Review the **Summary** tab information.

Option	Evaluation Process
Alert name and description	Review the name and description and verify that you are evaluating the alert for which you received an email message.
Recommendations	Review the top recommendation, and if available, other recommendations, to understand the steps that you must take to resolve the issue. If implemented, will the prioritized recommendations resolve the problem?
What is Causing the Issue?	Which symptoms were triggered? Which were not triggered? What affect does this evaluation have on your investigation? In this example, the alert that the datastore is running out of space is configured so that the criticality is symptom based. If you received a critical alert, then it is likely that the symptoms are already at a critical level, having moved up from Warning and Immediate. Look at the sparkline or metric graph chart for each symptom to determine when the problem escalated on the datastore object.

What to do next

- If you determine that the recommendations will resolve the problem, implement them. See Run a Recommendation on a Datastore to Resolve an Alert.
- If you need more information about the affected objects, continue your investigation. Begin by looking at other triggered symptoms for the data store. See Evaluate Other Triggered Symptoms for the Affected Data Store.

Evaluate Other Triggered Symptoms for the Affected Data Store

Because you need more information about the data store before you decide on the best response, you examine the **Symptoms** tab to see other triggered symptoms for the data store.

If other symptoms are triggered for the object besides the symptom included in the alert, evaluate them to determine what the symptoms reflect about the state of the object, and to decide whether the related recommendations might resolve the problem.

Prerequisites

Verify that you are addressing the alert for which you received an alert message in your email. See Respond to an Alert in Your Email.

Procedure

1 In the menu, click **Alerts** and select the alert name in the data grid.

The center pane view changes to display the alert detail tabs.

2 Click **View additional metrics > Alerts > Symptoms** and review the active symptoms.

Option	Evaluation Process
Criticality	Are other symptoms of similar criticality present that are affecting the object?
Symptom	Are any of the triggered symptoms related to the symptoms that triggered the current alert? Symptoms related to time remaining, capacity, or stress that could indicate storage problems?
Created On	Do the date and time stamps for the symptoms indicate that they were triggered before the alert you are investigating, indicating that it might be a related symptom? Were the symptoms triggered after the alert was generated, indicating that the alert symptoms contributed to these other symptoms?
Information	Can you identify a correlation between the alert symptoms and the other symptoms based on the triggering metric values?

What to do next

- If your review of the symptoms and the provided information clearly indicates that the recommendations will solve the problem, implement one or more of the recommendations. For an example of implementing one of the recommendations, see Run a Recommendation on a Datastore to Resolve an Alert.
- If your review of the symptoms did not convince you that the recommendations will resolve the problem or provide you with enough information to identify the root cause, continue your investigation using the **Events > Timeline** tab. See Compare Alerts and Events Over Time in Response to a Datastore Alert.

Compare Alerts and Events Over Time in Response to a Datastore Alert

To evaluate an alert over time, compare the current alert and symptoms to other alerts and symptoms, other events, other objects, and over time.

As a network operations engineer, you use the **Events > Timeline** tab to compare this alert to other alerts and events in your environment. This way, you can determine if you can resolve the problem of the datastore running out of disk space by applying one or more alert recommendations.

Prerequisites

Verify that you are addressing the alert for which you received an alert message in your email. See Respond to an Alert in Your Email.

Procedure

1 In the menu, click **Alerts** and select the alert name in the data grid.

The alert details appear to the right.

2 Click **View Events > Timeline**.

The **Timeline** tab displays the generated alert and the triggered symptoms for the affected object in a scrollable timeline format, starting when the alert was generated.

3 Scroll through the timeline using the week timeline at the bottom.

4 To view events that might contribute to the alert, click **Event Filters** and click the check box for each event type.

Events related to the object are added to the timeline. You add the events to your evaluation of the current state of the object and determine whether the recommendations can resolve the problem.

5 Click **View From** and select **Host** under Parents.

Because the alert is related to disk space, adding the host to the timeline enables you to see what alerts and symptoms are generated for the host. As you scroll through the timeline, ask: when did some of the related alerts begin? When are they no longer on the timeline? What was the effect on the state of the datastore object?

6 Click **View From** and select **Peer** under Parents.

If other datastores have alerts related to the alert you are currently investigating, seeing when the alerts for the other datastores were generated can help you determine what resource problems you are experiencing.

7 To remove canceled alerts from your timeline, click **Filters** and deselect the **Canceled** check box.

Removing the canceled alerts and symptoms from the timeline clears the view and enables you to focus on current alerts.

What to do next

- If your evaluation of alerts in the timeline indicated that one or more of the recommendations to resolve the alert are valid, implement the recommendations. See Run a Recommendation on a Datastore to Resolve an Alert.
- If you need more information about the affected object, continue your investigation. See View the Affected Datastore in Relation to Other Objects.

View the Affected Datastore in Relation to Other Objects

To view the object for which the alert was generated as it relates to other objects, use the topological map on the **Relationships** tab.

As a network operations engineer, you view a datastore and the related objects in a map to further your understanding of the problem. The map view helps determine if implementing the alert recommendations can resolve the problem.

Prerequisites

Evaluate the alert over time and in comparison to related objects. See Compare Alerts and Events Over Time in Response to a Datastore Alert.

Procedure

1. In the menu, click **Alerts**, select the alert name in the data grid, and click **View additional metrics > All Metrics**.
2. Click **Show Object Relationships**.

   The **Relationships** tab displays the datastore in a map with the related objects. By default, the badge that this alert affects is selected only on the toolbar. Objects in the tree show a colored square to indicate the current state of the badge.
3. To view the alert status of the objects for the other badges, click the **Health** button and then the **Efficiency** button.

   As you click each badge button, the squares on each object indicate whether an alert is generated and the criticality of the alert.
4. To view alerts for an object, select the object and click **Alerts**.

   The alert list dialog box appears, enabling you to search and sort for alerts for the object.
5. To view a list of the child objects for an object in the map, click the object.

   A list of the number of children by object type appears at the bottom of the center pane.
6. Use the options to evaluate the datastore.

   For example, what does the map tell you about the number of virtual machines that are associated with the datastore? If many virtual machines are associated with a datastore, moving them might free datastore disk space.

What to do next

- If your review of the map provided enough information to indicate that one or more of the recommendations to resolve the alert are valid, implement the recommendations. See Run a Recommendation on a Datastore to Resolve an Alert.
- If you need more information about the affected object, continue your investigation. See Construct Metric Charts to Investigate the Cause of the Data Store Alert.

Construct Metric Charts to Investigate the Cause of the Data Store Alert

To analyze the capacity metrics related to the generated alert, you create charts that compare different metrics. These comparisons help identify when something changed in your environment and what effect it had on the datastore.

As a network operations engineer, you create custom charts so that you can further investigate the problem, and to determine if implementing the alert recommendations will resolve the problem that the alert identifies.

Prerequisites

View the topological map for the data store to determine if related objects are contributing to the alert or if triggering symptoms indicate that the data store is contributing to other problems in your environment. See View the Affected Datastore in Relation to Other Objects.

Procedure

1. In the menu, click **Alerts**, select the alert name in the data grid, and click **View additional metrics > All Metrics**.

   The **Metric Charts** tab does not include charts. You must add the charts to compare.

2. To analyze the first recommendation, Add more capacity to the Datastore Storage, add related charts to the workspace.

   a. Enter **`capacity`** in the metric list search text box.

      The list displays metrics that contain the search term.

   b. Double-click the following metrics to add the following charts to the workspace:

      - Capacity | Used Space (GB)
      - Disk Space | Capacity (GB)
      - Summary | Number of Capacity Consumers

   c. Compare the charts.

      For example, if the Capacity | Used Space (%) chart shows an increase in used space, but the Disk Space | Capacity (GB) did not increase and the Summary | Number of Capacity Consumers did not decrease, then adding capacity is a solution, but it does not address the root cause.

3. To analyze the second recommendation, `vMotion some Virtual Machines to a different Datastore`, add related charts to the workspace.

   a. Enter **`vm`** in the metric list search text box.

   b. Double-click the **Summary | Total Number of VMs** metric to add it to the workspace

   c. Compare the 4 charts.

      For example, if the Summary | Total Number of VMs chart shows that the number of virtual machines did not increase enough to negatively affect the data store, then moving some of the virtual machines is a solution, but it does not address the root cause.

4 To analyze the third recommendation, Delete unused snapshots of virtual machines, add related charts to the workspace.

   a Enter **snapshot** in the metric list search text box.

   b Double-click the following metrics to add the charts to the workspace:

      - Disk Space | Snapshot Space (GB)
      - Disk Space Reclaimable | Snapshot Space | Waste Value (GB)

   c Compare the charts.

      For example, if the amount of Disk Space | Snapshot Space (GB) increased and the Disk Space Reclaimable | Snapshot Space | Waste Value (GB) indicates an area where space can be reclaimed, then deleting unused snapshots will positively affect the data store disk space problem and resolve the alert.

5 If this is a problematic data store that you must continue to monitor, you can create a dashboard.

   a Click the **Generate Dashboard** button on the workspace toolbar.

   b Enter a name for the dashboard and click **OK**.

      In this example, use a name like **Datastore disk space**.

   The dashboard is added to your available dashboards.

Results

You compared metric charts to determine if the recommendations are valid and which recommendation to implement first. In this example, the Delete unused snapshots of Virtual Machines recommendation appears to be the most likely way to resolve the alert.

What to do next

Implement the alert recommendations. See Run a Recommendation on a Datastore to Resolve an Alert.

Run a Recommendation on a Datastore to Resolve an Alert

As a network operations engineer, you investigated the alert regarding datastore disk space and determined that the provided recommendations can the problem. The recommendation to delete unused snapshots is especially useful. Use vRealize Operations Manager to delete the snapshots.

If you have not enabled actions in the vCenter adapter, you can manually delete the snapshots on your vCenter Server instance.

Prerequisites

- Compare the metric charts to identify the likely root cause of the alert. See Compare Alerts and Events Over Time in Response to a Datastore Alert .

Procedure

1. In the menu, click **Alerts** and select the alert name in the data grid. The alerts detail information appears on the right.
2. Review the Recommendations.

   Recommendations include the `Storage vMotion some virtual machines to a different datastore` recommendation and the `Delete unused snapshots for virtual machines` recommendation. The delete unused snapshot recommendation includes an action button.
3. Click **Delete Unused Snapshots for Datastore**.
4. In the **Days Old** text box, select or enter the number of days old the snapshot must be to be retrieved for deletions and click **OK**.

   For example, enter 30 to retrieve all snapshots on the datastore that are 30 days old or older.
5. In the **Delete Unused Snapshots for Datastore** dialog box, review the Snapshot Space, Snapshot Create Time, and the VM Name. Determine which snapshots to delete and select the check box for each one to delete.
6. Click **OK**.

   The dialog box that appears provides a link to Recent Tasks and a link to the task.
7. To verify that the task ran successfully, click **Recent Tasks**.

   The Recent Tasks page appears. The Delete Unused Snapshots action includes two tasks, one to retrieve the snapshots and one to delete the snapshots.
8. Select the Delete Unused Snapshot task that has the more recent finish time.

   This is the delete task. The status should be `Completed`.

Results

In this example, you ran an action on the datastore in vCenter Server. The other recommendations might also be valid.

What to do next

- Verify that the recommendations resolve the alert. Run a few collection cycles after you run the action and verify that the alert is canceled. Alerts are canceled when the conditions that generated them are no longer true.
- Implement the other recommendations. The other recommendations for this alert require you to use other applications. You cannot implement the recommendations from vRealize Operations Manager.

User Scenario: You See Problems as You Monitor the State of Your Objects

As you investigate your objects in the context of this scenario, vRealize Operations Manager provides details to help you resolve the problems. You analyze the state of your environment, examine current problems, investigate solutions, and take action to resolve the problems.

As a virtual infrastructure administrator, you regularly browse through vRealize Operations Manager at various levels so that you know the general state of the objects in your managed environment. Although no one has called or complained, and you do not see any new alerts, you are starting to see that your cluster is running out of capacity.

This scenario refers to objects that are associated with the VMware vSphere Solution, which connects vRealize Operations Manager to one or more vCenter Server instances. The objects in your environment include multiple vCenter Server instances, data centers, clusters (cluster compute resources), host systems, resource pools, and virtual machines.

As you perform the steps in this scenario, and progress through the stages of troubleshooting, you learn how to use vRealize Operations Manager to help you resolve problems. You will analyze the state of the objects in your environment, examine current problems, investigate solutions, and take action to resolve the problems.

This scenario shows you how to evaluate the problems that occur on your objects, and take action to resolve problems.

- With the Analysis tab, you view the settings for object resources, click the links provided to further analyze the problem, and examine the policy settings and thresholds.
- Using the Events tab, you examine the symptoms that triggered on the objects, determine when the problems that triggered those symptoms occurred, identify the events associated with those problems, and examine the metric values involved.
- On the Details tab, you investigate the metric activity as a graph, list, or distribution chart, and view the heat maps to examine the criticality levels of your objects.
- With the Environment tab, you evaluate the health, risk, and efficiency of various objects as they relate to your overall object hierarchy. You view the object relationships to determine how an object that is in a critical state might be affecting other objects.

To support future troubleshooting and ongoing maintenance, you can create a new alert definition, and create a dashboard and one or more views and reports. To plan for growth and account for newly approved projects, you can create and commit capacity projects. To enforce the rules used to monitor your objects, you can create and customize operational policies.

Prerequisites

Verify that you are monitoring one or more vCenter Server instances.

Procedure

1 Analyze the State of Your Environment

The Analysis tabs help you analyze your objects in multiple ways. As a Virtual Infrastructure Administrator, you use the Analysis tabs to evaluate the details about the state of your objects to help you resolve problems.

2 Troubleshoot Problems with a Host System

You use the Troubleshooting tabs to identify the root cause of problems that are not resolved by alert recommendations or simple analysis.

3 Examine the Environment Details

Examine the status of your objects in the views and heatmaps so that you can identify the trends and spikes that are occurring with the resources on your cluster and objects. To determine whether any deviations have occurred, you can display overall summaries for an object, such as for the cluster disk space usage breakdown.

4 Examine the Environment Relationships

You use the Environment Overview and List to examine the status of the badges as they relate to the objects in your environment hierarchy, and determine which objects are in a critical state for a particular badge. To view the relationships between your objects to determine whether an ancestor object that has a critical problem might be causing problems with the descendants of the object, you use the Environment Map.

5 Fix the Problem

You use the analysis and troubleshooting features of vRealize Operations Manager to examine problems that put your objects in a critical state, and identify solutions. To resolve the problems, where actions exist for the object type, you select an object and an available action that is specific to the object. Or, you can open the object in the vSphere Web Client and modify the object settings to resolve the problem.

6 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.

7 Create Dashboards and Views

To help you investigate and troubleshoot problems with your cluster and host systems that might occur in the future, you can create dashboards and views that apply the troubleshooting tools and solutions that you used to research and solve the problems with your host system, to make those troubleshooting tools and solutions available for future use.

Analyze the State of Your Environment

The Analysis tabs help you analyze your objects in multiple ways. As a Virtual Infrastructure Administrator, you use the Analysis tabs to evaluate the details about the state of your objects to help you resolve problems.

As you browse through the inventory tree, you notice that one of your clusters, named USA-Cluster, is experiencing capacity problems. You use the Analysis tabs to begin to investigate the cause of the problem on USA-Cluster, and you start to see problems reported with the capacity on one of your host systems and other objects.

Prerequisites

Verify that you understand the context of this scenario. See User Scenario: You See Problems as You Monitor the State of Your Objects.

Procedure

1. In the menu, click **Environment**, then in the left pane click **vSphere Hosts and Clusters**and select the object.
2. Click the **Analysis** tab.

   You see red icons on the Capacity Remaining and Time Remaining tabs.
3. Click the **Time Remaining** tab.

   You see that the memory allocation is severely constrained.
4. View the time remaining breakdown for the cluster.

   The icons indicate that zero days remain, with no planned capacity projects considered.
5. Scroll down until you see the Time Remaining in Related Objects pane.

   The parent object is the data center, and the peer represents another cluster. The child objects include the resource pool and host systems. The data center and one of the host systems are experiencing critical memory problems.
6. Hover your mouse over the red parent and child icons.

   The memory capacity has expired on the data center and one of the host systems.

Results

The memory capacity problem on the cluster is affecting the memory capacity of the related objects.

What to do next

Use the Troubleshooting tab to further troubleshoot the capacity problems on your cluster and host system.

Troubleshoot Problems with a Host System

You use the Troubleshooting tabs to identify the root cause of problems that are not resolved by alert recommendations or simple analysis.

To further troubleshoot the symptoms of the capacity problems that are occurring on the cluster and host system, and determine when those problems occurred, you use the Troubleshooting tabs to continue to investigate the memory problem.

Prerequisites

Use the Analysis tabs to analyze your environment. See Analyze the State of Your Environment.

Procedure

1. In the menu, click **Environment**, then in the left pane click **vSphere Hosts and Clusters** and select the object. For example, USA-Cluster.
2. Click the **Alerts** tab and review the symptoms.

   The **Symptoms** tab displays the symptoms that triggered on the selected cluster. You notice that several critical symptoms exist.

   - `Cluster Compute Resource Time Remaining with committed projects is critically low`
   - `Cluster Compute Resource Time Remaining is critically low`
   - `Capacity remaining is critically low`
3. Analyze the critical symptoms.
   a. Hover your mouse over each critical symptom to identify the metric used.
   b. To view only the symptoms that affect the cluster, enter **`cluster`** in the quick filter text box.

      When you hover over `Cluster Compute Resource Time Remaining is critically low`, the metric `Badge|Time Remaining with committed projects (%)` appears. You notice that its value is less than or equal to zero, which caused the capacity symptom to trigger and generate an alert on USA-Cluster.
4. Click the **Events > Timeline** tab to review the triggered symptoms, alerts, and events that occurred on USA-Cluster over time, and identify when the problems occurred.
   a. Click the calendar and select **Last 7 Days** as the range.

      Several events appear in red.
   b. Hover your mouse over each event to view the details.
   c. To display the events that occurred on the cluster's data center, click **View From**, and select **Datacenter**.

      Warning events for the data center appear in yellow.

d Hover your mouse over the warning events.

You notice that the density is starting to get low, and that a hard threshold violation occurred on the data center late in the evening. The hard threshold violation shows that the Badge|Density metric value was under the acceptable value of 25, and that the violation triggered with a value of 14.89.

e To view the affected child objects, click **View From** and select **Host System**.

5 Click the **Events** tab to examine the changes that occurred on USA-Cluster, and determine whether a change occurred that contributed to the root cause of the alert or other problems with the cluster.

a Review the graph.

By reviewing the graph, you can determine whether a reoccurring event has caused the errors. Each event indicates that the guest file system is out of disk space. The affected objects appear in the pane below the graph.

b Click each red triangle to identify the affected object and highlight it in the pane below.

6 Click the **All Metrics** tab to evaluate the objects in their context in the environment topology to help identify the possible cause of a problem.

a In the top view, select **USA-Cluster**.

b In the metrics pane, expand **Badge** and double-click **Badge|Capacity Remaining (%)**.

The Badge|Capacity Remaining (%) calculation is added to the lower right pane.

c In the metrics pane, double-click **Density**.

d In the metrics pane, double-click **Workload**.

e On the toolbar, click **Date Controls** and select **Last 7 Days**.

The metric chart indicates that the capacity for the cluster remained at a steady level for the past week, but that the cluster density increased to its maximum value in the last several days. The Badge|Workload (%) calculation displays the workload extremes that correspond to the density problem.

Results

You have analyzed the symptoms, timeline, events, and metrics related to the problems on your cluster, and determined that the heavy workload on the cluster has decreased the cluster density in the last several days, which indicates that the cluster is starting to run out of capacity.

What to do next

Examine the Details views and heatmaps to interpret the properties, metrics, and alerts to look for trends and spikes that occur in the resources for your objects, the distributions of resources across your objects, and data maps to examine the use of various resource types across your objects.

Examine the Environment Details

Examine the status of your objects in the views and heatmaps so that you can identify the trends and spikes that are occurring with the resources on your cluster and objects. To determine whether any deviations have occurred, you can display overall summaries for an object, such as for the cluster disk space usage breakdown.

To examine the problems with your USA-Cluster further, use the Details views to display the metrics and collected capacity data for your cluster. Each view includes specific metrics data collected from your objects. For example, trend views use data collected from objects over time to generate trends and forecasts for resources such as memory, CPU, disk space, and so on.

Use the heatmaps to examine the capacity levels on the cluster, host systems, and virtual machines. The block sizes and colors are based on the metrics selected in the heatmap configuration. For example, the heatmap that shows the most abnormal workload for virtual machines is sized by the Badge|Workload (%) metric, and is colored by the Badge|Anomaly metric.

Prerequisites

Use the Troubleshooting tabs to look for root causes. See Troubleshoot Problems with a Host System.

Procedure

1. Click **Environment > vSphere Hosts and Clusters > USA-Cluster**.
2. Examine the detailed information about USA-Cluster in the views.
   a. Click the **Details** tab and click **Views**.

      The views provide multiple ways to look at different types of collected data by using trends, lists, distributions, and summaries.

   b. In the search text box, enter **`capacity`**.

      The list filters and displays the capacity views for clusters and other objects.

   c. Click the view named **Cluster Capacity Risk Forecast**, and examine the number of virtual machines for USA-Cluster in the lower pane.

      Even though the USA-Cluster has two host systems and 30 virtual machines, no capacity exists.
3. Examine the host systems in the cluster, and reclaim capacity from the descendant virtual machines.
   a. Click the **Analysis** tab, and click **Capacity Remaining**.

   b. In the inventory tree, expand **USA-Cluster**, and click each of the host systems.

      The host system named w2-vcopsqe2-009 is in a critical state, with no capacity remaining.

c In the lower pane, expand **Memory**, and expand **Allocation**.

The stress free value is zero, and the amount of memory available is zero, which indicates that the capacity of the host system has been depleted.

d Click the **Details** tab, and click **Views**, and click the **Virtual Machine Reclaimable Capacity** view.

e In the lower pane, click the title of the **Reclaimable Memory** column to sort the list of virtual machines so that the largest amount of reclaimable capacity is on top.

f To reclaim capacity from several virtual machines, click to the right of the first virtual machine name, then press **Shift** and click to the right of the last virtual machine that has capacity to reclaim.

The virtual machines that have reclaimable capacity are highlighted.

g Click the gear icon, and select **Set CPU Count and Memory for VM**.

h Click the **Current CPU** column title to sort the list according to the highest number of CPUs.

Based on the actual use of the virtual machines listed, the **New CPU** column recommends fewer CPUs for each virtual machine.

i Click the check box next to each virtual machine that has a recommended lower CPU count, and click **OK**.

By reducing the number of CPUs for each virtual machine, you free up capacity on your host system, and improve the USA-Cluster capacity and workload.

4 Examine the heatmaps for the host system and virtual machine objects in USA-Cluster.

a In the inventory tree, click **USA-Cluster**.

b Click **Details**, click **Heatmaps**, and click through the list of heatmap views.

c Click **Which VMs currently have the highest CPU demand and contention?**

The heatmap displays blocks that represent the objects in USA-Cluster. The block for a virtual machine appears in red, which indicates that it has a critical problem.

d Hover over the red block and examine the details.

The cluster, host system, and virtual machine names appear, with links to more information about the object.

e Click **Show Sparkline** to display the activity trend on the virtual machine.

f Click each of the **Details** links to display more information.

Results

To verify that freeing up memory on the virtual machines has improved the workload of the host system and the cluster, you can now examine the status of the host system and cluster.

You used views and heatmaps to evaluate the status of your objects and identify trends and spikes, and free up capacity for your host system and USA-Cluster. To further narrow in on problems, you can examine the other views and heatmaps. You can also create your own views and heatmaps.

What to do next

Examine the badge status for the objects in your environment hierarchy to determine which objects are in a critical state, and examine the object relationships to determine whether a problem on one object is affecting one or more other objects.

Examine the Environment Relationships

You use the Environment Overview and List to examine the status of the badges as they relate to the objects in your environment hierarchy, and determine which objects are in a critical state for a particular badge. To view the relationships between your objects to determine whether an ancestor object that has a critical problem might be causing problems with the descendants of the object, you use the Environment Map.

As you click each of the badges in the Environment Overview, you see that several objects are experiencing critical problems with health, workload, and faults. Others are reporting critical risk status, and many are in critical time remaining and capacity remaining states.

Several objects are experiencing stress. You notice that you can reclaim capacity from multiple virtual machines and a host system, but the overall efficiency status for your environment displays no problems.

Prerequisites

Examine the status of your objects in views and heatmaps. See Examine the Environment Details.

Procedure

1 Click **Environment > vSphere Hosts and Clusters > USA-Cluster**.

2 Examine the USA-Cluster environment overview to evaluate the badge states of the objects in a hierarchical view.

a In the inventory tree, click **USA-Cluster**, and click **Environment > Overview**.

b On the Badge toolbar, click through the badges and look for red icons to identify critical problems.

Option	Evaluation Process
Status icons	When the status of my object is critical, what must I do to resolve the problem? How can I be notified before serious problems occur?
Badges: Health, Workload, Anomalies, and Faults	How might the health and workload of my host systems be affecting my virtual machines? Are anomalies and faults on my host systems and virtual machines affecting other objects?
Badges: Risk, Time Remaining, Capacity Remaining, Stress	How does the stress level of my cluster and host systems affect the virtual machines descendants?
Badges: Efficiency, Reclaimable Capacity, Density	To improve efficiency, how can I reclaim capacity from the cluster, host systems, resource pool, and virtual machines, and apply the reclaimed capacity to other objects in my environment?

As you click through the badges, you notice that your vCenter Server and other top level objects appear to be healthy, but you see that a host system and several virtual machines are in a critical state for health, workload, and faults. Several objects also have critical problems with time remaining and capacity remaining.

c Hover your mouse over the red icon for the host system to display the IP address.

d Enter the IP address in the search text box, and click the link that appears.

The host system is highlighted in the inventory tree. You can then look for recommendations or alerts for the host system on the Summary tab.

3 Examine the environment list and view the badge status for your objects to determine which objects are in a critical state.

a Click **Environment > List**.

b Examine the badge states for the objects in USA-Cluster.

c Click the **Capacity Remaining** badge column name to sort the object list and display the objects that are in a critical state.

Many of the objects that are at risk for capacity remaining also display critical states for time remaining, risk, and health. You notice that multiple virtual machines and a host system named w2-vropsqe2-009 are critically affected. Because the host system is experiencing the most critical problems, and is likely affecting other objects, you must focus on resolving the problems with the host system.

d Click the host system named **w2-vropsqe2-009**, which is in a critical state, to locate it in the inventory tree.

e Click **w2-vropsqe2-009** in the inventory tree, and click the **Summary** tab to look for recommendations and alerts so that you can take action.

4 Examine the environment map.

a Click **Environment > Map**.

b In the inventory tree, click **USA-Cluster**, and view the map of related objects.

In the relationship map, you can see that the USA-Cluster has an ancestor data center, one descendant resource pool, and two descendant host systems.

c Click the host system named **w2-vropsqe2-009**.

The types and numbers of descendant objects for this host system appear in the list below. Use the descendant object list identify all of the objects related objects to the host system that might be experiencing problems.

What to do next

Take action in the user interface to resolve the problems.

Fix the Problem

You use the analysis and troubleshooting features of vRealize Operations Manager to examine problems that put your objects in a critical state, and identify solutions. To resolve the problems, where actions exist for the object type, you select an object and an available action that is specific to the object. Or, you can open the object in the vSphere Web Client and modify the object settings to resolve the problem.

You have used the Analysis, Troubleshooting, Details, and Environment areas of the user interface to examine the critical problems that occur on your objects. To resolve those problems, you can select actions from the Actions menu, which appears in list and view menus, and various dashboard widgets.

The actions that you can select are specific to an object type, such as a virtual machine. Although you can select an action when you have selected a host system that is experiencing critical problems related to capacity and time, all but one of the actions that you can take apply to virtual machines. The action to delete unused snapshots applies to datastores.

Prerequisites

Examine the environment relationships. See Examine the Environment Relationships.

Procedure

1 In the menu, click **Environment**, then click **vSphere Hosts and Clusters > vSphere World** in the left pane.

2 From the **Details** view, select the host system and take action.

a In the inventory tree, click the host system named **w2-vropsqe2-009**.

b Click **Details > Views**, and enter **memory** in the search text box.

c Click the view named **Host Rightsizing CPU, Memory, and Disk Space**.

The host system named w2-vropsqe2-009 appears in the lower pane. You see that the provisioned CPUs and memory for the host system are wasting capacity, and realize that you can free up some capacity in an attempt to resolve the capacity problem on the host system.

Provisioned	Recommendation	Reclaimable
16 Core CPUs	10 Core CPUs	35 Core CPUs
127 GB memory	35 GB memory	68 GB memory
4,011 GB disk space	11,158 GB disk space	122 GB disk space

d In the lower pane, click to the right of the host system named **w2-vropsqe2-009**.

e On the toolbar in the lower pane, click the **Open in external application** icon, and click **Open Host in vSphere Client**.

f Log in to the vSphere Web Client, and modify the provisioned CPU and memory for the host system.

3 (Optional) From the Environment view, select the host system and take action.

a In the inventory tree, click **USA-Cluster**.

b Click **Environment > List**.

c Click to the right of the name of the w2-vropsqe2-009 host system.

d In the lower pane, click to the right of the host system named **w2-vropsqe2-009**.

e On the toolbar in the lower pane, click the **Open in external application** icon, and click **Open Host in vSphere Client**.

f Log in to the vSphere Web Client, and modify the provisioned CPU and memory for the host system.

4 (Optional) From the inventory tree, select the host system and take action.

a In the inventory tree, click **w2-vropsqe2-009**.

b At the top of the toolbar in the right pane, click **Actions**.

c Click **Open Host in vSphere Client**.

d Log in to the vSphere Web Client, and modify the provisioned CPU and memory for the host system.

Results

You have used the available actions to resolve problems on a host system that is experiencing critical problems. The available action appears in **Content > Actions**.

What to do next

To become aware of critical problems on your objects before they adversely affect the performance of other objects and your environment, create an alert definition, and optionally add actions to the alert definition recommendations.

Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for vRealize Operations Manager to alert you. When the alert is triggered on your host system, vRealize Operations Manager alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have vRealize Operations Manager notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

Procedure

1. In the menu, click **Alerts** and then in the left pane, select **Alert Settings > Alert Definitions**.
2. Enter **`capacity`** in the search text box.

   Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.
3. Click the plus sign to create a new capacity alert definition for your host systems.
   a. In the alert definition workspace, for the Name and Description, enter **`Hosts – Alert on Capacity Exceeded`**.
   b. For the Base Object Type, select **vCenter Adapter > Host System**
   c. For the Alert Impact, select the following options.

Option	Selection
Impact	Select **Risk**.
Criticality	Select **Immediate**.
Alert Type and Subtype	Select **Application : Capacity**.
Wait Cycle	Select **1**.
Cancel Cycle	Select **1**.

d For Add Symptom Definitions, select the following options.

Option	Selection
Defined On	Select **Self**.
Symptom Definition Type	Select **Metric / Supermetric**.
Quick filter (Name)	Enter **capacity**.

e From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the right pane.

In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.

f For Add Recommendations, enter **virtual machine** in the quick filter text box.

g Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the right pane.

This recommendation is set to Priority 1.

4 Click **Save** to save the alert definition.

Your new alert appears in the list of alert definitions.

Results

You have added an alert definition to have vRealize Operations Manager alert you when the capacity of your host systems begins to run out.

Create Dashboards and Views

To help you investigate and troubleshoot problems with your cluster and host systems that might occur in the future, you can create dashboards and views that apply the troubleshooting tools and solutions that you used to research and solve the problems with your host system, to make those troubleshooting tools and solutions available for future use.

To readily view the status of your cluster and host systems when your CIO asks you about their health, you can use the decision support dashboards on the vRealize Operations Manager Home page. For example, you can:

- Use the vSphere Clusters dashboard to view the utilization index, CPU demand, and memory use for your clusters. This dashboard also tracks the net use and disk I/O operations.
- Use vSphere Cluster Configuration Summary dashboard to track the high availability status, and other configuration items.
- Use the vSphere Hosts Overview to examine the capacity levels of your cluster, host systems, and virtual machines.
- Use the Health of Host Systems dashboard to view the active alert list, capacity metric chart and heatmap for your host system.

Or, you might need to create your own dashboards to track the status of your clusters and host systems.

If you work in a Network Operations Center environment and have multiple monitors, you can run multiple instances of vRealize Operations Manager, and dedicate a monitor to each specific dashboard so that you can visually track the status of your objects.

Prerequisites

Create an alert definition to alert you when the capacity of your host system is getting low. See Create a New Alert Definition.

Procedure

1. In the menu, click **Dashboards** and look through the list of existing dashboards to determine whether you can use the cluster and host system dashboards to track your clusters and host systems.
2. Click the **Health of Host Systems** dashboard, and review the widgets included on it.

   The inclusion of the Object List, Alert List, Metric Picker, Metric Chart, Heatmap, and Top-N widgets would allow you to easily peruse the status of the host systems that you select in the Object List widget. This dashboard has the widget interaction configured so that the object you select in the Object List widget is the object for which the other widgets display data.
3. Create and configure a new dashboard that has widgets to monitor the health of your host systems and generate alerts.
   a. Above the dashboard view, click **Actions** and select **Create Dashboard**.
   b. In the New Dashboard workspace, for the Dashboard Name, enter **`Health of Host Systems`**, and leave the other default settings.
   c. In the Widget List workspace, add the Object List widget and configure it to display host system objects.
   d. Add the Alert List widget to the dashboard, and configure it to display capacity alerts when the capacity of your host systems becomes an immediate risk.
   e. In the Widget Interactions workspace, for each widget listed, select the Object List widget as the provider to drive the data to the other widgets, and click **Apply Interactions**.
   f. In the Dashboard Navigation workspace, select the dashboards that receive data from the selected widgets, and click **Apply Navigations**.

   After vRealize Operations Manager collects data, if a problem occurs with the capacity of your host systems, the Alert List widget on your new dashboard displays the alerts that are configured for your host systems.

What to do next

Prepare to share information with others, plan for growth and new projects, and use policies to continuously monitor all of the objects in your environment. See Reports, Chapter 6 Planning the Capacity for Your Managed Environment Using vRealize Operations Manager, and Policies .

Monitoring and Responding to Alerts

Alerts indicate a problem in your environment. Alerts are generated 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 is generated, you are presented with the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

Alerts notify you when an object or group of objects are exhibiting symptoms that are unfavorable for your environment. By monitoring and responding to alerts, you stay aware of problems and can react to them in a timely fashion.

Generated alerts drive the status of the top level badges, Health, Risk, and Efficiency.

In addition to responding to alerts, you can generally respond to the status of badges for objects in your environment.

You cannot assign alerts to vRealize Operations Manager users. Your users must take ownership of an alert.

Monitoring Alerts in vRealize Operations Manager

You can monitor your environment for generated alerts in several areas in vRealize Operations Manager. The alerts are generated when the symptoms in the alert definition are triggered, letting you know when the objects in your environment are not operating within the parameters you defined as acceptable.

Generated alerts appear in many areas of vRealize Operations Manager so that you can monitor and respond to problems in your environment.

Alerts

Alerts are classified as Health, Risk, or Efficiency. Health alerts indicate problems that require immediate attention. Risk alerts indicate problems that must be addressed in the near future, before the problems become immediate health problems. Efficiency alerts indicate areas where you can reclaim wasted space or improve the performance of objects in your environment.

You can monitor the alerts for your environment in the following locations.

- Alerts
- Health
- Risk
- Efficiency

You can monitor alerts for a selected object in the following locations.

- Alert Details, including the **Summary**, **Timeline**, and **Metric Charts** tabs
- **Summary** tab
- **Alerts** tab
- **Events** tab
- Custom dashboards
- Alert notifications

Working with Alerts

Alerts indicate a problems that must be resolved so that triggering conditions no longer exist and the alert is canceled. Suggested resolutions are provided as recommendations so that you can approach the problem with solutions.

As you monitor alerts, you can take ownership, suspend, or manually cancel alerts.

When you cancel an alert, the alert and any symptoms of type fault, message event, or metric event are canceled. You cannot manually cancel other types of symptoms. If the alert was triggered by a fault symptom, message event symptom or metric event symptom, then the alert is effectively canceled. If the alert was triggered by a metric symptom or property symptom, a new alert might be created for the same conditions in the next few minutes.

The correct way to remove an alert is to address the underlying conditions that triggered the symptoms and generated the alert.

Migrated Alerts

If you migrated alerts from a previous version of vRealize Operations Manager, the alerts are listed in the overview with a cancelled status, but alert details are not available.

User Scenario: Monitor and Process Alerts in vRealize Operations Manager

Alerts in vRealize Operations Manager notify you when objects in your environment have a problem. This scenario illustrates one way that you can monitor and process alerts for the objects for which you are responsible.

An alert is generated when one or more of the alert symptoms are triggered. Depending on how the alert is configured, the alert is generated when one symptom is triggered or when all of the symptoms are triggered.

As the alerts are generated, you must process the alerts based on the negative affect they have on objects in your environment. To do this, you start with Health alerts, and process them based on criticality.

As a virtual infrastructure administrator, you review the alerts at least twice a day. As part of your evaluation process in this scenario, you encounter the following alerts:

- Virtual machine has unexpected high CPU workload

- Host has memory contention that a few virtual machines cause
- Cluster has many virtual machines that have memory contention because of memory compression, ballooning, or swapping

Procedure

1 In the menu, click **Alerts**.

2 Select **Time** in the Group By filter and the click the down arrow in the Created On column, so the most recent alerts are listed first .

3 In All Filters, select **Criticality > Warning**

You have listed all the Warning alerts in order of when they fired, with the most recent alerts appearing first.

4 Review the alerts by name, the object on which it was triggered, the object type, and the time at which the alert was generated.

For example, do you recognize any of the objects as objects that you are responsible for managing? Do you know that the fix that you will implement in the next hour will fix any of the alerts that are affecting the Health status of the object? Do you know that some of your alerts cannot be resolved at this time because of resource constraints?

5 To indicate to other administrators or engineers that you are taking ownership of the `Virtual machine has unexpected high CPU workload` alerts, click the selected alerts, click **Actions** on the menu bar, and click **Take Ownership**.

The Assigned to: field in Alert Details updates with your user name. You can only take ownership of alerts, you cannot assign them to other users.

6 To take ownership and temporarily exclude the alert from affecting the state of the object, select the `Host has memory contention caused by a few virtual machines` alert in the list, click **Actions** on the menu bar, and click **Suspend**.

a Enter `60` to suspend the alert of an hour.

b Click **OK**.

The alert is suspended for 60 minutes and you are listed as the owner in the alert list. If it is not resolved in an hour, it returns to an active state.

7 Select the row that contains the `Cluster has many Virtual Machines that have memory contention due to memory compression, ballooning or swapping` alert, click **Actions** on the menu bar, and click **Cancel Alert** to remove the alert from the list.

This alert is a known problem that you cannot resolve until the new hardware arrives.

The alert is removed from the alert list, but the underlying condition is not resolved by this action. The symptoms in this alert are based on metrics, so the alert will be generated during the next collection and analysis cycle. This pattern continues until you resolve the underlying hardware and workload distribution issues.

Results

You processed the critical health alerts and took ownership of the ones to resolve or troubleshoot further.

What to do next

Respond to an alert. See User Scenario: Respond to a vRealize Operations Manager Alert in the Health Alert List.

User Scenario: Respond to a vRealize Operations Manager Alert in the Health Alert List

Generated alerts in vRealize Operations Manager appear in the alert lists. You use the alert lists to investigate, resolve, and begin troubleshooting problems in your environment.

In this scenario, you investigate and resolve the `Virtual machine has unexpected high CPU workload` alert. The alert might be generated for more than one virtual machine.

Prerequisites

- Process and take ownership of the alerts you will troubleshoot and resolve. See User Scenario: Monitor and Process Alerts in vRealize Operations Manager.
- Review information about how the Power Off Allowed setting works when you run actions. See Working With Actions That Use Power Off Allowed.

Procedure

1. In the menu, click **Alerts**.
2. To limit the list to virtual machine alerts, click **All Filters** on the toolbar.
   a. Select **Object Type** in the drop-down menu.
   b. Enter **`virtual machine`** in the text box.
   c. Click **Enter**.

   The alerts list displays only alerts based on virtual machines.
3. To locate the alerts by name, enter **`high CPU workload`** in the **Quick filter (Alert)** text box.
4. In the list, click the **Virtual machine has unexpected high CPU workload** alert name.
5. Review the information. Click **Alert Settings > Recommendations** in the left pane to show the recommendations.

Option	Evaluation Process
Alert Description	Review the description so that you better understand the alert.
Recommendations	Do you think that implementing one or more of the recommendations will resolve the alert?

Option	Evaluation Process
What is Causing the Issue?	Do the triggered symptoms support the recommendations? Do the other triggered symptoms contradict the recommendation, indicating that you must investigate further? In this example, the triggered symptoms indicate that the virtual machine CPU demand is at a critical level and that the virtual machine anomaly is starting to get high.
Non-Triggered Symptoms	Some alerts are generated only when all the symptoms are triggered. Others are configured to generate an alert when any one of the symptoms are triggered. If you have non-triggered symptoms, evaluate them in the context of the triggered alerts. Do the non-triggered symptoms support the recommendations? Do the non-triggered symptoms indicate that recommendations are not valid and that you must investigate further?

6 To resolve the alert based on the recommendation to check the guest applications to determine whether high CPU workload is an expected behavior, click the **Action** menu on the center pane toolbar and select **Open Virtual Machine in vSphere Client**.

   a Log in to the vCenter Server instance using your vSphere credentials.

   b Launch the console for the virtual machine and identify which guest applications are consuming CPU resources.

7 To resolve the alert based on the recommendation to add more CPU capacity to this virtual machine, click **Set CPU Count for VM**.

   a Enter a new value in the **New CPU** text box.

   The value that appears is the calculated recommended size. If vRealize Operations Manager was monitoring the virtual machine for six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.

   b Select the following options to allow power off or to create a snapshot, depending on how your virtual machines are configured.

Option	Description
Power Off Allowed	Shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system. In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied.
Snapshot	Creates a snapshot of the virtual machine before you add CPUs. If the CPU is changed with CPU Hot Plug enabled, then the snapshot is taken with the virtual machine running, which consumes more disk space.

   c Click **OK**.

The action adds the recommended number of CPUs to the target virtual machine.

8 Allow several collection cycles to run after implementing the recommended changes and check the alert list.

What to do next

If the alert does not reappear after several collection cycles, it is resolved. If it reappears, further troubleshooting is required. For an alternative scenario for troubleshooting alerts, see User Scenario: An Alert Arrives in Your Inbox.

Monitoring and Responding to Problems

The organization of the tabs and options in vRealize Operations Manager provides a built-in workflow that you can use when you work with objects in your environment.

The tabs, **Summary**, **Alerts**, **Analysis**, and so on, provide a progressive level of detail about the selected object. As you work through the tabs, starting with the high level **Summary** and **Alerts** tabs, you see the general state of an object. If you identify a problem, you use the aggregated metrics in the **Analysis** tabs to view the state of the object in a more detail. The data provided in the **Events** tabs is useful when you are investigating the root cause of a problem. The **Details** tabs are specific data views and the **Environment** tabs show object relationships.

As you monitor objects in your environment, you will discover which tabs provide the information that you need when you are investigating problems.

Evaluating Object Information Using Badge Alerts and the Summary Tab

The Summary tab that is associated with the other object tabs summarizes Heath, Risk, and Efficiency badge alerts for the selected object and displays the top alerts that lead to the current state.

Use this tab as an overview of alerts for an object, object group, or application - to evaluate the affect that alerts are having on an object and to begin troubleshooting problems. For more detail on the badge Alerts, click **Badge Alerts**, further to the right on the tool bar.

Badge Alert Types

The Health, Risk, and Efficiency badge states are based on the number and criticality of the generated alerts for the selected object.

- Health alerts indicate problems that affect the health of your environment and require immediate attention to ensure that service to your customers is not affected.
- Risk alerts indicate problems that are not immediate threats but should be addressed in the near future.
- Efficiency alerts tell you where you can improve performance or reclaim resources.

Alerts for an Object or an Object Group

When you are working with a single object, the Top alerts are the alerts generated for the object and the Top Alerts for Children are the alerts generated for any child or other descendant objects in the currently selected navigation hierarchy. For example, if you are working with a host object in the vSphere Host and Clusters navigation hierarchy, children can include virtual machines and datastores.

When you are working with object groups, which can include one object type, such as hosts, or multiple objects types, such as hosts, virtual machines, and datastores, all the group member objects are children of the group container. The most critical generated alerts for the member objects appear as Top Alerts for Children.

For an object group, the only Top Alerts that might be generated are the predefined group population alerts. A group population alert considers the health of all group members and is triggered if the average health is above the Warning, Immediate, or Critical threshold. If a group population alert is generated, then the badge score and color is affected by the alert. If a group population alert is not generated, then the badges are green. This behavior is because an object group is a container for other objects.

Summary Tab and Related Hierarchies

The alerts that appear on the **Summary** tab for an object can vary depending on the currently selected hierarchy in the Related Hierarchies in the left pane.

Depending on the selected hierarchy, you see different alerts and relationships on the **Summary** tab for an object. The current focus object name is on the center pane title bar, but the children alerts depend on the relationships that the highlighted hierarchy defined in the Related Hierarchies list in the upper left pane. For example, if you are working with a host object relative to virtual machines in the vSphere Hosts and Clusters hierarchy, then children commonly include virtual machines and datastores. But if you are working with the same host as a member of an object group, then any alerts on virtual machines that are also members of the group do not appear because the host and the virtual machines are considered children of the group and peers among each other. In this example, the focus of the **Summary** tab is the host in the context of the group, not the vSphere Hosts and Clusters hierarchy.

Summary Tab Evaluation Techniques

You can evaluate the state of objects, starting with the **Summary** tab, by using one or more of the following techniques.

- Select an object or object group, click on the alerts on the **Summary** tab, and resolve the problems that the alert indicates.
- Select an object and examine the information about the current object that is provided in the other tabs. For example, you start on the object **Summary** tab and compare the generated alerts to the analytic information about the object on the **Analysis** tabs.
- Select an object, review the alerts on the **Summary > Alerts** tab, and select other objects, comparing the volume and types of alerts generated for different objects.

User Scenario: Evaluate the Badge Alerts for Objects for a vRealize Operations Manager Object Group

In vRealize Operations Manager, you use alerts on a group to review the summary alert information for hosts and virtual machine descendant objects so that you can see how the state of one object type can affect the state of the other.

As a network operations center engineer, you are responsible for monitoring a group of hosts and virtual machines for the sales department. As part of your daily tasks, you check the state of the objects in the group to determine if there are any immediate problems or any upcoming problems based on generated alerts. To do this you start with your group of objects, particularly the host systems in the group, and review the information in the **Summary** tab.

In this example, the group includes the following object alerts.

- `Host has memory contention caused by a few virtual machines` is a Health alert
- `Virtual Machine has chronic high memory workload` is a Risk alert
- `Virtual Machine is demanding more CPU than the configured limit` is a Risk alert
- `Virtual Machine has large disk snapshots` is an Efficiency alert

The following method of evaluating alerts on the **Summary** tab is provided as an example for using vRealize Operations Manager and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

Prerequisites

- Create a group that includes virtual machines and the hosts on which they run. For example, Sales Dept VMs and Hosts. For an example of how to create a similar group, see Create a Custom Accounting Department Group.
- Review how the **Summary** tab works with object groups and related hierarchies. See Evaluating Object Information Using Badge Alerts and the Summary Tab.

Procedure

1. In the menu, click **Environment** .
2. Click the **Custom Groups** tab and click, for example, your **Sales Dept VMs and Hosts** group.
3. To view the alerts for a host and the associated child virtual machines, in the left pane, click, for example, **Host System** and click the host name in the lower left pane.

   The **Summary** tab displays the Health, Risk, and Efficiency badges, the top alerts for the host. (Because the group is still the focus, the alerts for the child virtual machines do not appear in the Top Alerts for Descendants widgets at the **Badge Alerts** tab.)
4. To view the Summary tab for the host so that you can also work with the child virtual machines, click the right arrow to the right of the host name in the lower left pane.

5 Select the **vSphere Hosts and Clusters**, located in the upper part of the left pane.

To work with alerts for child virtual machines, the host in the vSphere Hosts and Clusters hierarchy must be the focus of the **Summary** tab rather than the host as member of the object group.

6 To view the alert details for an alert in the list, click the alert name.

When multiple objects are affected, and you click the alert link to view the details, the Health Issues dialog box appears. If there is only one object affected, the **Alerts** tab for the object is displayed.

7 On the **Alerts** tab, begin evaluating the recommendations and triggered symptoms.

In this scenario, a recommendation for this generated alert is to move some virtual machines with high memory workload from this host to a host with more available memory.

8 To return to the object **Summary** tab so that you can review alerts for any child virtual machines, click the back button located in the left pane.

The host is again the focus of the object **Summary** tab. Generated alerts for the child virtual machines appear below.

9 Click on each virtual machine alert and evaluate the information provided on the **Alerts** tab.

Virtual Machine Alert	Evaluation
Virtual Machine has chronic high memory workload	The recommendation is to add more memory to this virtual machine. If one or more virtual machines are experiencing high workload, this situation is probably contributing to the host memory contention alert. These virtual machines are candidates for moving to a host with more available memory. Moving the virtual machines can resolve the host memory contention alert and the virtual machine alert.
Virtual Machine is demanding more CPU than the configured limit	The recommendations include increasing or removing the CPU limits on this virtual machine. If one or more virtual machines are demanding more CPU than is configured, and the host is experiencing memory contention, then you cannot add CPU resources to the virtual machine without further stressing the host. These virtual machines are candidates for moving to a host with more available memory. Moving the virtual machines would allow you to increase the CPU count and resolve the virtual machine alert, and might resolve the host memory contention alert.

10 Based on your evaluation, take action based on the child virtual machine recommendations.

Results

After you take action, it will take a few collection cycles to determine if your actions resolved the virtual machine and host alerts.

What to do next

After a few collection cycles, look again at your Sales VMs and Hosts group to determine if the alerts are canceled and no longer appear in the object **Summary** tab. If the alerts are still present, see User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options for an example troubleshooting workflow.

Summary Tab

The Summary tab provides an overview of the state of the selected object, group, or application. Use this tab to evaluate the impact that alerts are having on the object and use the information to begin troubleshooting problems.

How the Summary Tab Works

Based on the object selected, the following summary tabs are displayed:

- VM Summary Tab
- Datastore Summary Tab
- Host Summary Tab
- Cluster Summary Tab
- Custom Group and Container Summary Tab

Where to Find the Summary Tab

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object.
- In the menu, select **Alerts** to display the All Alerts screen. Click an **alert** to bring up the alert details on the right, then click **View Additional Metrics** to see more information about the alert and the object that triggered the alert. Click the **Summary** tab.

Understanding the Summary Tab

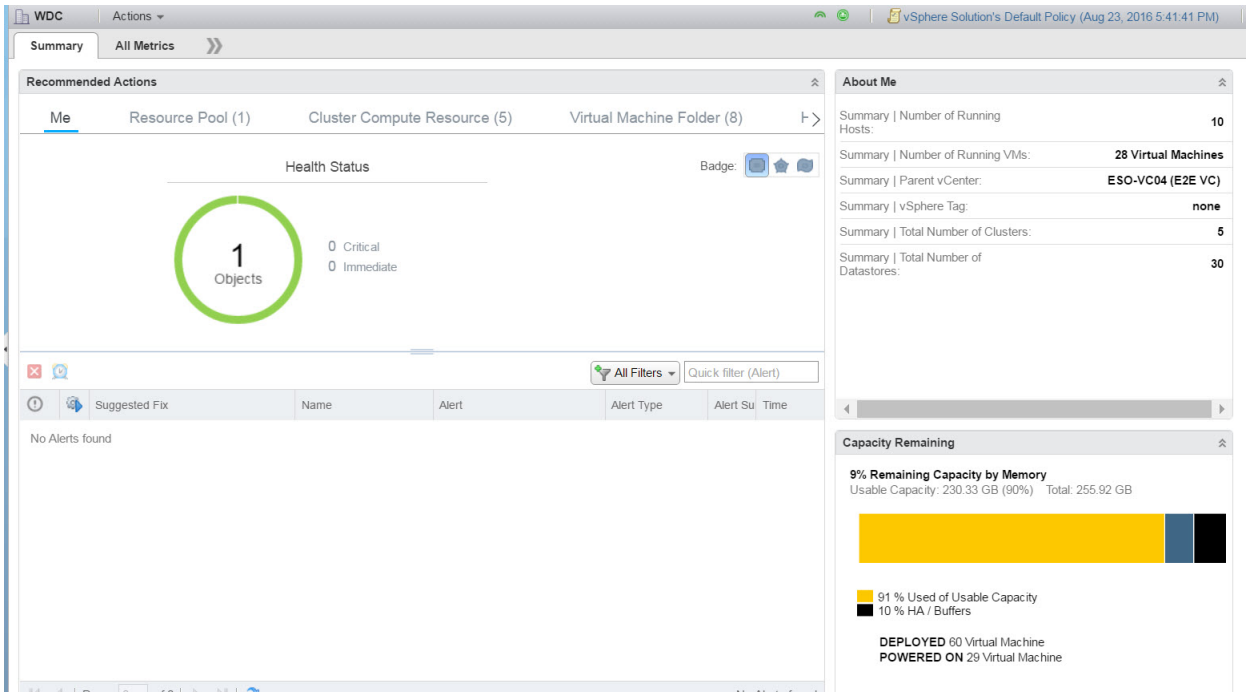

Table 5-1. Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems in the near future ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.
About Me	This widget displays the summary of metrics and properties of the selected object for review.
Capacity Remaining	This widget displays a score indicating the remaining computing resources as a percent of the total consumer capacity for the most constrained resource.

Datastore Summary Tab

The Datastore Summary tab provides an overview of the state of the selected datastore. For the selected object, the Datastore Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the datastore and use the information to begin troubleshooting problems.

Understanding the Datastore Summary Tab

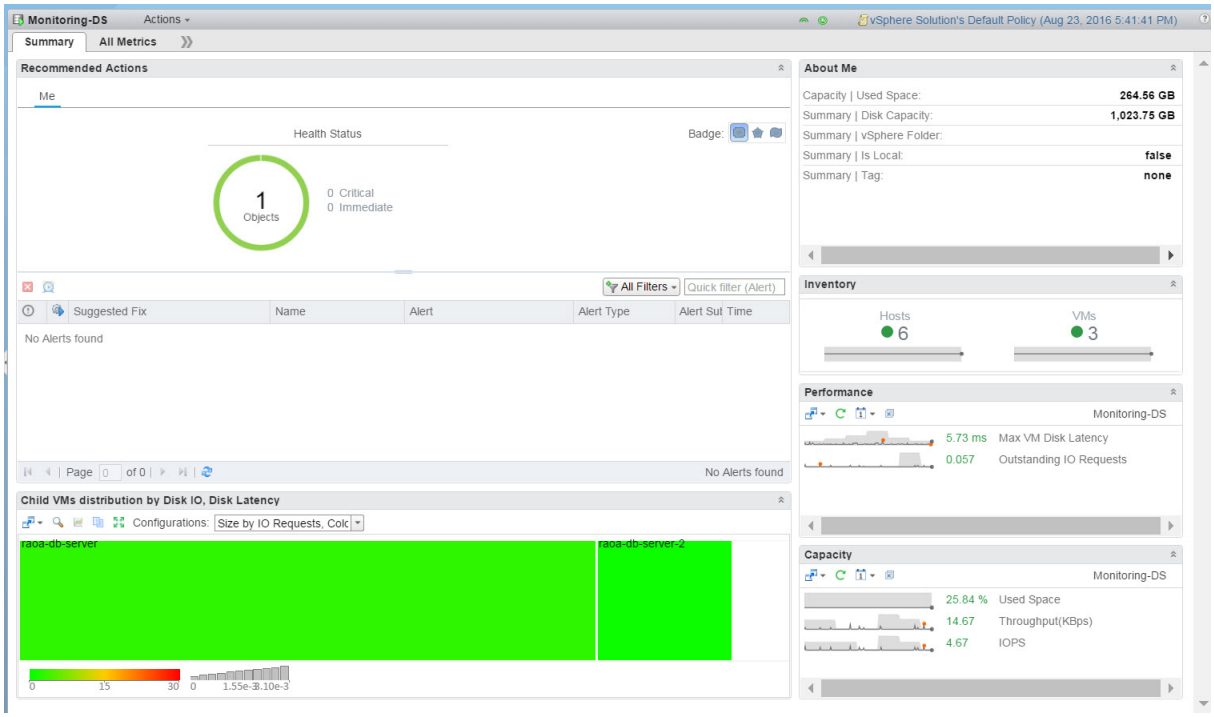

Table 5-2. Datastore Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems in the near future ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.
About Me	This widget displays the key metrics and properties of the selected object.
Inventory	This widget displays the number of hosts and VMs associated with the datastore.
Capacity	This widget displays a visual summary of the capacity and workload resources used by the objects in your environment. It displays the latest value and a trend line of the various key indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the expanded chart.
Performance	This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the expanded chart.
Child VMs distribution by Disk IO, Disk Latency	As per the configuration that you choose from the list, this widget displays heat maps to show the distribution of the child VMs based on the Disk IO and Disk latency metrics. It helps to quickly evaluate the status of all the VMs using the same datastore. It also helps to check if there are problems that impact all the VMs or if a group of VMs is the source of a problem.

Host Summary Tab

The Host Summary tab provides an overview of the state of the selected host. For the selected object, the Host Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the host and use the information to begin troubleshooting problems.

Understanding the Host Summary Tab

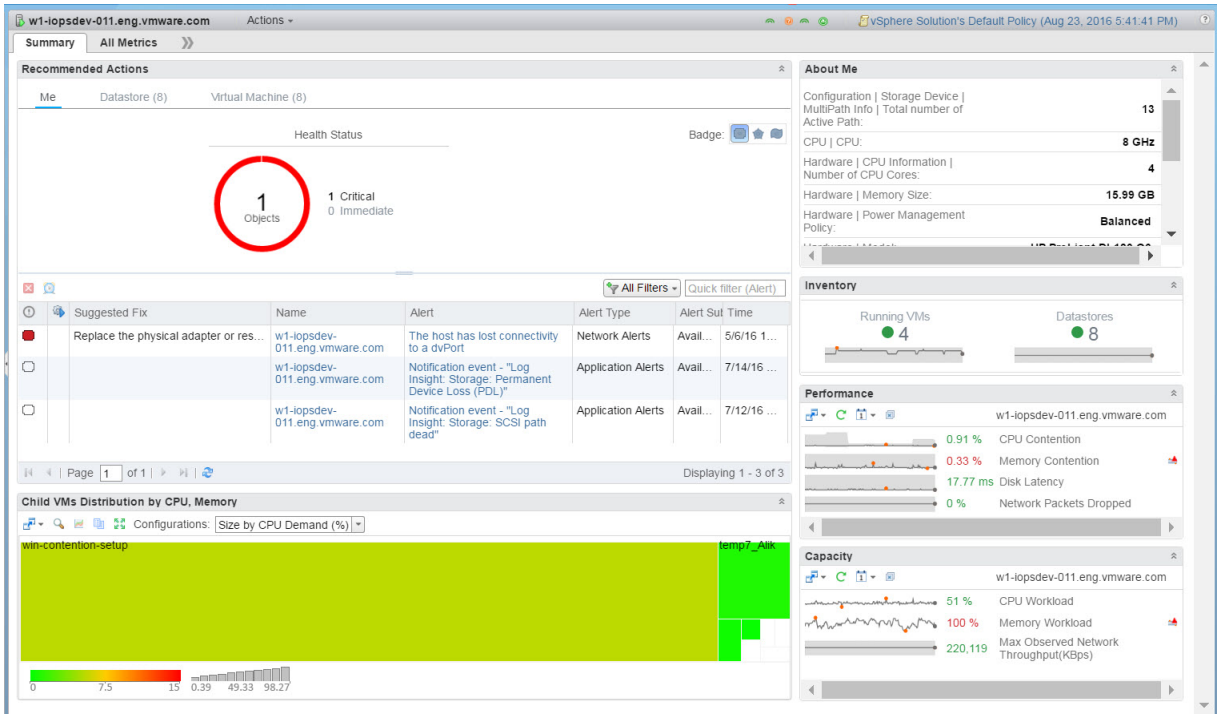

Table 5-3. Host Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems soon ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.
About Me	This widget displays the key metrics and properties of the selected object.
Inventory	This widget displays the number of running VMs and Datastores associated with the selected host.
Capacity	This widget displays a visual summary of the capacity and workload resources used by the objects in your environment. It displays the latest value and a trend line of the various key indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the detailed chart.

Table 5-3. Host Summary Tab Options (continued)

Option	Description
Performance	This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the expanded chart.
Child VMs Distribution by CPU, Memory	As per the configuration that you select from the list, this widget displays the heat maps showing the distribution of the child VMs based on CPU and Memory metrics. It also helps to identify the noisy VMs in the host.

VM Summary Tab

The VM Summary tab provides an overview of the state of the selected VM. For the selected object, the VM Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the VM and use the information to begin troubleshooting problems.

Understanding the VM Summary Tab

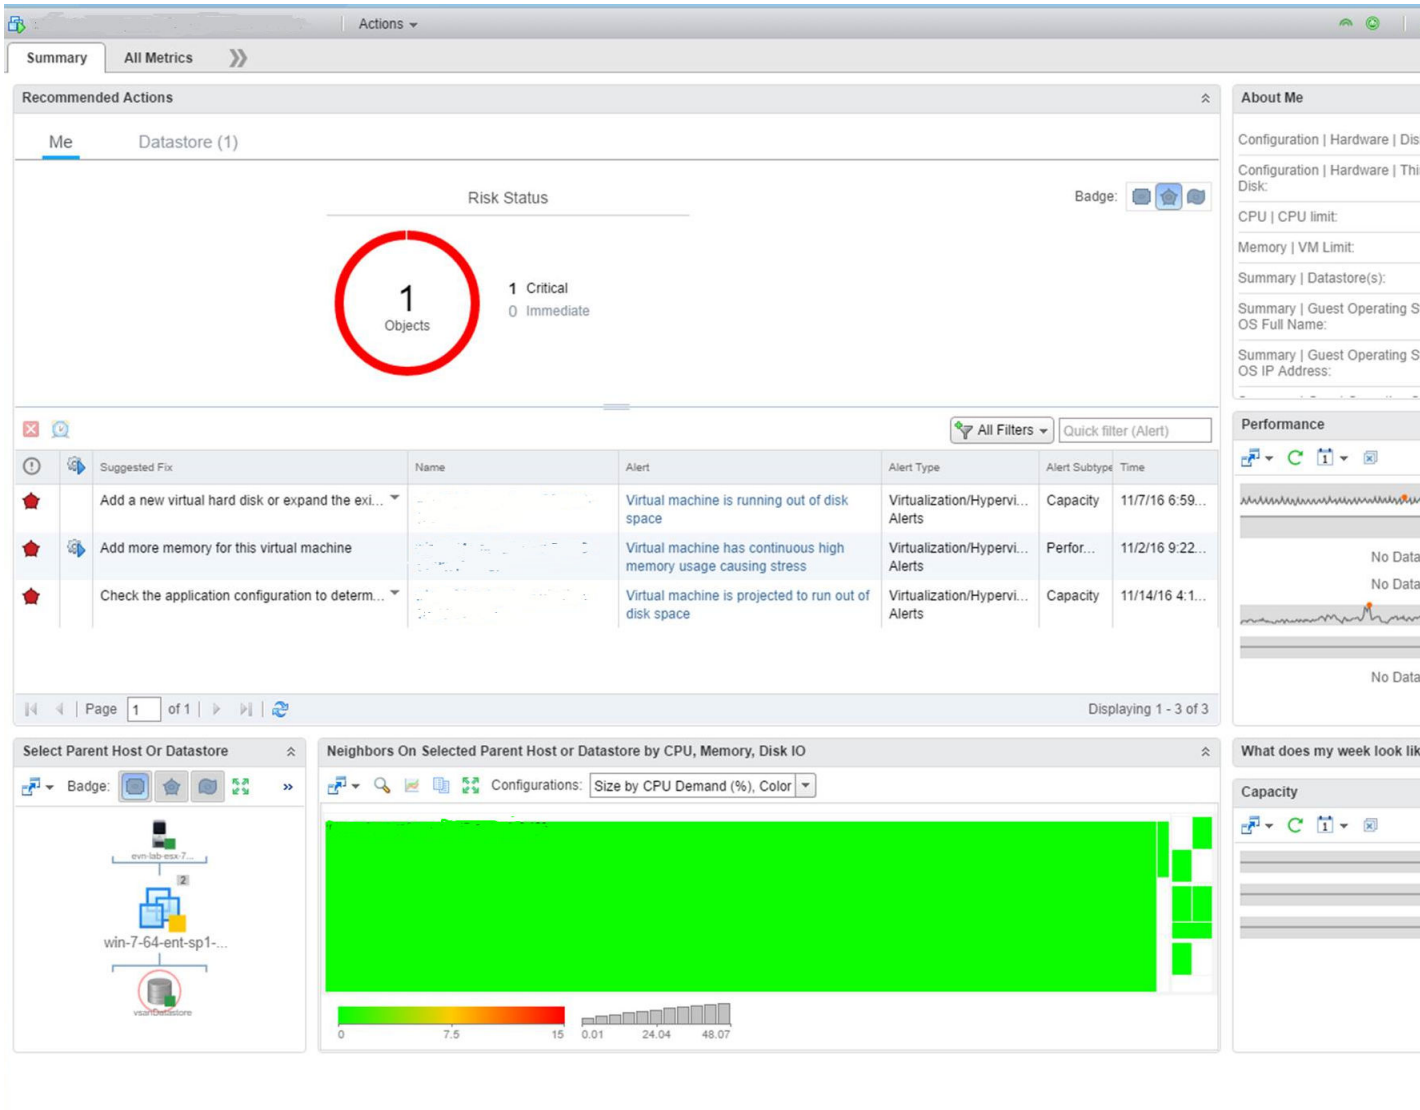

Table 5-4. VM Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems in the near future ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.
About Me	This widget displays the key metrics and properties of the selected object.

Table 5-4. VM Summary Tab Options (continued)

Option	Description
Capacity	This widget displays a visual summary of the capacity and workload resources used by the objects in your environment. It displays the latest value and a trend line of the various key indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the expanded chart.
Performance	This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Double click each metric to see the expanded chart.
What does my week look like?	This widget displays a quick view of the amount of stress that the VM went through in the last week per day. It also helps to identify the pattern of load on the VM during the week.
Select Parent Host or Datastore	This widget display the status of the parent host or datastore of the selected VM. This input controls the data displayed in the heat map.
Neighbors on Selected Parent Host or Datastore by CPU, Memory, Disk IO	As per the configuration that you choose from the list, this widget displays heat maps showing the distribution of the neighbors on selected parent host or datastore by CPU, Memory, and Disk IO. It helps to identify the noisy neighbours using the same infrastructure.

Cluster Summary Tab

The Cluster Summary tab provides an overview of the state of the selected cluster. For the selected object, the Cluster Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the cluster and use the information to begin troubleshooting problems.

Understanding the Cluster Summary Tab

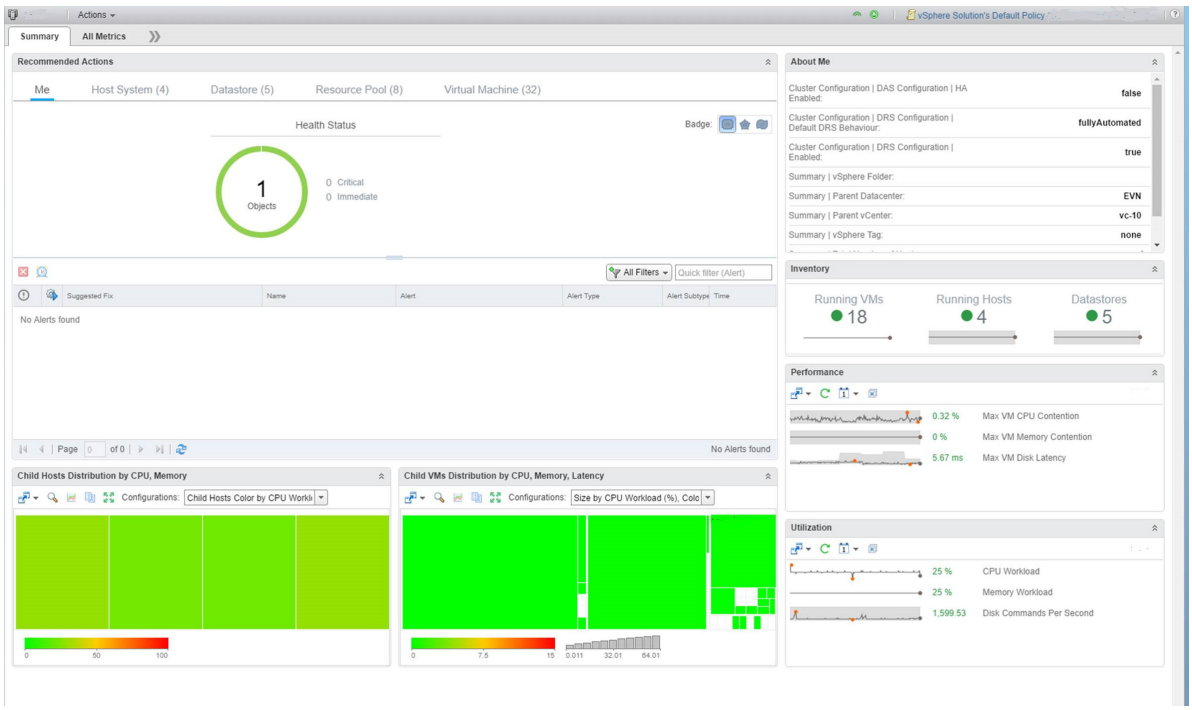

Table 5-5. Cluster Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems in the near future ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.
About Me	This widget displays the key metrics and properties of the selected object.
Inventory	This widget displays the number of running hosts, running VMs, and datastores associated with the cluster.
Utilization	This widget provides a summary of the utilization of the cluster by CPU/Memory and IO. It displays a trend line for the last 24 hours and the latest value in the color associated with its health based on the symptom associated with this metric.
Performance	This widget displays the trend line of maximum KPI values for any of the VMs running on the cluster for the last 24 hours. It also displays the latest value in a colour that represents its health based on the symptom associated with this metric. Click each metric to see a detailed view of the chart.

Table 5-5. Cluster Summary Tab Options (continued)

Option	Description
Child Hosts Distribution by CPU, Memory	As per the configuration that you choose from the list, the heat map shows the distribution of the child hosts based on CPU and memory. It helps to quickly identify VMs with high demand and VMs with latency problems.
Child VMs Distribution by CPU, Memory, Latency	As per the configuration that you choose from the list, the heat map shows the distribution of the child VMs based on CPU, memory, and latency. This heat map helps to identify hosts with high workloads.

Custom Group and Container Summary Tab

The Custom Group and Container Summary tab provides an overview of the state of the selected group or a container. For the selected object, the Custom Group and Container Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the group or a container and use the information to troubleshoot the problems.

Understanding the Custom Group and Container Summary Tab

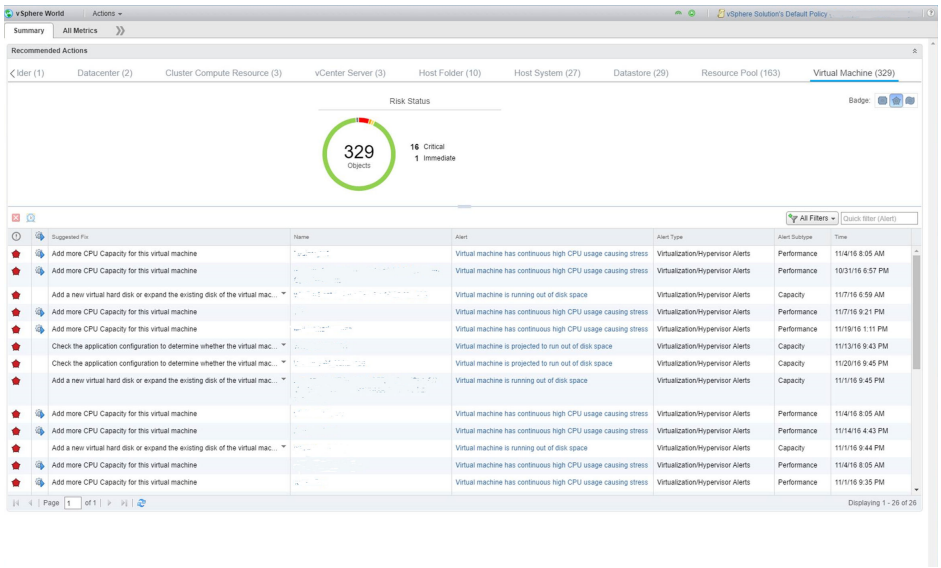

Table 5-6. Customer Group and Container Summary Tab Options

Option	Description
Recommended Actions	This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance. The badges provide a visual indicator of the alert status for the following alert types. ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you should look into any problems in the near future ■ Efficiency alerts indicating that you can reclaim resources. Click the badge to see the alerts for the object.

Investigating Object Alerts

The **Alerts** tab provides a list of generated alerts for the currently selected object. When you are working with objects, reviewing and responding to generated alerts on the **Alert** tab helps you manage problems in your environment.

The alerts notify you when a problem occurs in your environment based on configured alert definitions. Object alerts are useful to you as an investigative tool in two ways. They can provide you with proactive notification about problems in your environment before a user calls you to complain, and they provide information about the object that you can use when troubleshooting general or reported problems.

As you review the **Alerts** tab, you can add ancestors and descendants to the list to broaden your view of the alerts. You can see if alerts on the current object affect other objects or how the current object is affected by the problems indicated by alerts on other objects.

Depending on the best practices and workflows of your infrastructure operations team, you can use the object **Alerts** tab to manage generated alerts on individual objects.

- Take ownership of alerts so that your team knows that you are working to resolve the problem.
- Suspend an alert so that is temporarily excluded from affecting the Health, Risk, or Efficiency state of the object while you investigate the problem.
- Cancel alerts that you know are a result of a deliberate action, for example, a network card was removed from a host for replacement, or that are known issues that you cannot resolve at this time because of resource constraints. Canceling an alert that is generated because of only fault, message event, or metric event symptoms cancels the alert permanently. Canceling an alert that is generated because of metric, super metric, or property symptoms can result in the alert being regenerated if the underlying metric or property condition remains true. It is only effective to cancel alerts generated because of fault, message event, or metric event symptoms.

Investigating and resolving alerts helps you provide the best possible environment to your customers.

User Scenario: Respond to Alerts on the Alerts Tab for Problem Virtual Machines

You respond to alerts for objects so that you can bring the affected objects back to the required level of configuration or performance. Based on the information in the alert and using other information provided in vRealize Operations Manager, you evaluate the alert, identify the most likely solution, and resolve the problem.

As a virtual infrastructure administrator or operations manager, you troubleshoot problems with objects. Reviewing and responding to the generated alerts for objects is part of any troubleshooting process. In this example, you want to resolve workload problems for a virtual machine. As part of that process, you review the **Alerts** tab to determine what alerts might indicate or contribute to the identified problem.

The problem virtual machine is db-01-kyoto, which you use as a database server.

The following method of responding to alerts is provided as an example for using vRealize Operations Manager and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

Prerequisites

- Verify that the vCenter Adapter has been configured for the actions in each vCenter Server instance.
- Verify that you understand how to use the Power Off Allowed option if you are running Set CPU Count, Set Memory, and Set CPU Count and Memory actions. See Working With Actions That Use Power Off Allowed.

Procedure

1. Enter the name of the object, **`db-01-kyoto`**, in the **Search** text box and select the virtual machine in the list.

   The object **Summary** tab appears. The Top Alerts panes display important active alerts for the object.

2. Click the **Analysis** tab.

   The **Workload** tab is the first tab. This badge indicates that the workload is highest by CPU, but memory is also above the configured limit.

3. Click the **Alerts** tab.

   In this example, the alert list includes the follow alerts that might be related to the problem you are investigating.

   - Virtual machine has unexpected high CPU workload.
   - Virtual machine has unexpected high memory workload.

4 In the upper left pane, select the **vSphere Hosts and Clusters** related hierarchy and select ancestor or descendant alerts to add to the list.

You want to check for possible alerts on ancestor or descendant objects in the context of the selected hierarchy.

a On the toolbar, click **Show Ancestor Alerts** and select the **Host System** and **Resource Pool** check boxes.

Any alerts for the host system or resource pool related to this virtual machine are added to the list.

b Click **Show Descendant Alerts** and select **Datastore**.

Any alerts for the datastore are added to the list.

In this example, there are no additional alerts for the host, resource pool, or datastore, so you begin addressing the virtual machine alerts.

5 Click the **Virtual machine has unexpected high CPU workload** alert name.

The **Alert Details Summary** tab appears.

6 Review the recommendations to determine if one or more suggested recommendations can fix the problem.

This example includes the following common recommendations:

- `Check the guest applications to determine whether high CPU workload is expected behavior.`
- `Add more CPU capacity for this virtual machine.`

7 To follow the `Check the guest applications to determine whether high CPU workload is expected behavior` recommendation, click **Actions** on the title bar and select **Open Virtual Machine in vSphere Client**.

The vSphere Web Client Summary tab appears so that you can open the virtual machine in the console and check which applications are are contributing to the reported high CPU workload.

8 To follow the `Add more CPU Capacity for this virtual machine` recommendation, click **Set CPU Count for VM** .

a Enter a value in the **New CPU** text box.

The default value that appears before you provide a value is a recommended value based on analytics.

b To allow the action to power off the virtual machine before running the action if Hot Add for CPU is not enabled, select the **Power Off Allowed** check box.

c To create a snapshot before changing the virtual machine CPU configuration, select the **Snapshot** check box.

d Click **OK**.

e Click the Task ID link and verify that the task ran successfully.

The specified number of CPUs are added to the virtual machine.

What to do next

After a few collection cycles, return to the object **Alerts** tab. If the alert no longer appears, then your actions resolved the alert. If the problem is not resolved, see User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options for an example troubleshooting workflow.

Alerts Tab

The Alerts tab is a list of all the alerts generated for the selected object, group, or application. Use the alerts list to evaluate the number of generated alerts for the object so that you can begin resolving them.

How the Alerts Tab Works

All the active alerts for the selected object appear in the list. By default, the alerts are grouped by Time. You can select multiple rows in the list using Shift+click, Control+click. Modify the filter if you want to see inactive alerts.

Manage the alerts in the list using the toolbar options. Click the **alert name** to see the alert details for the affected object. The alert details appear on the right, including the symptoms triggered along with the alert. The system offers recommendations for addressing the alert and links to additional information. A **Run Action** button may appear in the details. Hover over the button to learn what recommendation is performed if you click the button. Click the **X** at the top right of the alert details to return to the list view.

To see the object details, click the **Summary** Tab.

Where You Find the Alerts Tab

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. Click the **object** to bring up the object's **Summary** tab. Click the **Alerts > Alerts** tabs.
- In the menu, select **Search** and locate the object of interest. Click the **object** to bring up the object's **Summary** tab. Click the **Alerts > Alerts** tabs.

Alerts Tab 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. Additional toolbar options enable you to review parent and child alerts related to the alert you are reviewing. Use the data grid to view the alerts and alert details.

Table 5-7. Actions Menu

Option	Description
Actions menu	Select an alert from the list to turn on the Actions menu, then select an option from the menu.
Menu Options:	
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.
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.
Go to Alert Definition	Switches to the Alert Definitions page, with the definition for the previously selected alert displayed.
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 5-8. View From Menu

Options	Description
Self	The selected object.
Parents <options>	Displays the alerts for the ancestors of the selected object. Parents in this instance include the parents, grandparents, and so on, of the object. For example, the parents of a host are a folder, storage pod, cluster, data center, and vCenter Server instance.
Children <options>	Displays the alerts for the descendants of the selected object. Children in this instance include the children and grandchildren of the object. For example, the descendants of a host are datastores, resources pools, and virtual machines.

Table 5-9. 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 "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 5-10. Alerts 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	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.
Created On	Date and time when the alert was generated.
Status	Current state of the alert. Possible values include Active or Canceled.

Table 5-10. Alerts Data Grid (continued)

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.

Table 5-11. 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 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.
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.
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.

Table 5-11. All Filters (continued)

All Filters	Descriptions
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.

Evaluating Metric Information

The **All Metrics** tab provides a relationship map and user-defined metric charts. The topological map helps you evaluate objects in the context of their place in your environment topology. The metric charts are based on the metrics for the selected object that you think helps identify the possible cause of a problem in your environment.

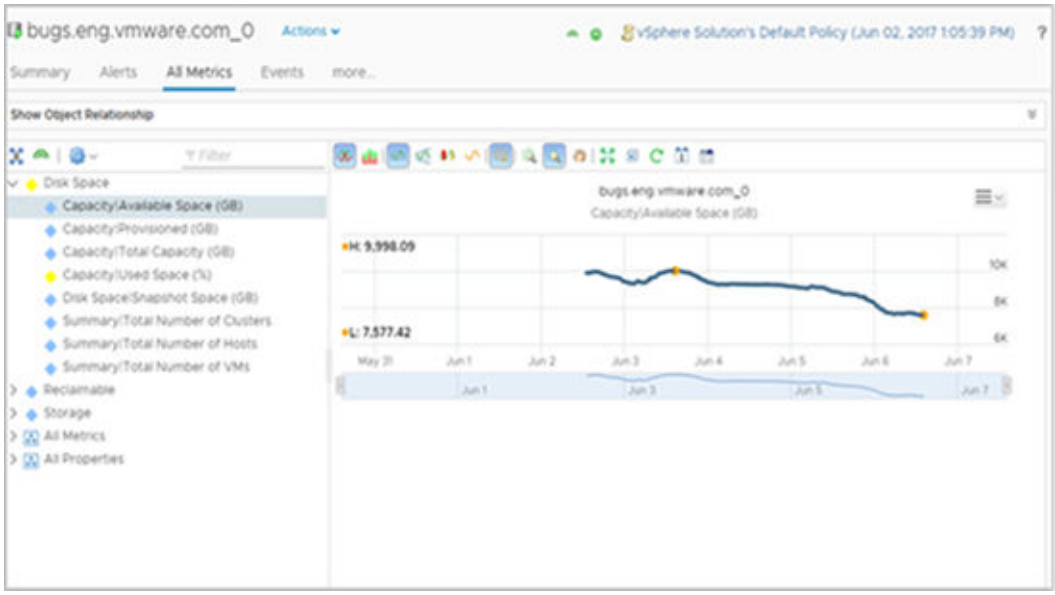

Although you might be investigating problems with a single object, for example, a host system, the relationship map allows you to see the host in the context of parent and child objects. It also works as a hierarchical navigation system. If you double-click an object in the map, that object becomes the focus of the map. The available metrics for the object become active in the lower-left pane.

You can also build your own set of metric charts. You select the objects and metrics that provide you with a detailed view of changes to different metrics for a single object, or for related objects over time.

Where available, the **All Metrics** tab provides pre-defined sets of metrics to help you when looking at a specific aspect of an object. For example, if you have a problem with a host, access the most relevant information about the host by looking at the metrics displayed in the pre-defined lists. You can edit these groups of metrics, and create additional groups, by dragging and dropping metrics and properties from the All Metrics and All Properties lists.

For more information about the metrics, refer to Chapter 7 Metric, Property, and Alert Definitions.

Where You Find the All Metrics Tab

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the objects that you want.

Create Metric Charts When You Troubleshoot a Virtual Machine Problem

You create a custom group of metric charts when you troubleshoot a problem with a virtual machine so that you can compare different metrics. The level of detail that you can create using the **All Metrics** tab, can contribute significantly to your effort to find the root cause of a problem.

As an administrator investigating a performance problem with a virtual machine, you determined that you must see detailed charts about the following reported symptoms.

- Guest file system overall disk space usage reaching critical limit
- Guest partition disk space usage

The following method of evaluating problems using the **All Metrics** tab is provided as an example for using vRealize Operations Manager and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

Procedure

1. Enter the name of the virtual machine in the **Search** text box on the menu bar.

   In this example, the virtual machine name is **sales-10-dk**.
2. Click the **All Metrics** tab.
3. In the relationship topology map, click the virtual machine, **dk-new-10**.

   The metrics list, located in the left of the center pane, displays virtual machine metrics.

4 On the chart toolbar, click **Date Control** and select a time that is on or before the symptoms were triggered.

5 Add metric charts to the display area for the virtual machine.

   a In the metric list, select **Guest Files System Stats > Total Guest File System Free (GB)** and double-click the metric name.

   b To add the guest partition, for example, C:\, select **Guest Files System Stats > C:\ > Guest File System Free (GB)** and double-click the metric name.

   c To add disk space for comparison, select **Disk Space > Capacity Remaining (%)** and double-click the metric name.

6 Compare the charts.

   You can see a decrease in the file system free space, and that the virtual machine disk space capacity remaining is decreasing at a steady rate. You determine that you must add disk space to the virtual machine. However, you do not know if the datastore can support the change to the virtual machine.

7 Add the datastore capacity chart to the charts.

   a In the topology map, double-click the host.

      The topology map refreshes with the host as the focus object.

   b Click the datastore.

   c In the metric list, which is updated to display datastore metrics, select **Capacity > Available Space (GB)** and double-click the metric name.

8 To determine if sufficient capacity is available on the datastore to support increasing the disk space on the virtual machine, review the datastore capacity chart.

Results

You know that you must increase the size of the virtual disk on the virtual machine.

What to do next

Expand the virtual disk on the virtual machine and assign it to stressed partitions. Click **Actions**, on the object title bar, and view the virtual machine in the vSphere Web Client.

Troubleshooting with the All Metrics Tab

The **All Metrics** tab provides a relationship map and metric charts. The topological map helps you evaluate objects in the context of their place in your environment topology. Metric charts are based on the metrics for the active map object that you think can help you identify the cause of a problem.

How All Metrics Works

The relationships map shows the selected object, the related objects, and the number of generated alerts for each one. If you double-click an object icon, the selected object becomes the focus of the map. The map is updated for the selected object, and the metrics list shows only the metrics for the selected object.

Using the metrics list, you create charts based on metrics that you think can help you investigate problems. You customize the charts to evaluate the data in more detail. To save the configured charts, you create a dashboard using the toolbar option.

Where available, the metrics list also displays pre-defined groups of metrics that contain the most relevant metrics for the selected object. You can edit these groups, and create your own customized groups of metrics by dragging and dropping metrics and properties from the All Metrics and All Properties lists.

Where You Find All Metrics

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the objects that you want.

All Metrics Options

The options include the map toolbar, the metric selector options, the metric charts toolbar, and the toolbar on each chart.

Table 5-12. Relationship Map

Option	Description
Badge	Displays the state of the selected badge on each object in the map.
Zoom to fit	Resizes the map to fit in the available space.
Pan	Click and drag the map so that you can view a particular object in the map regardless of the level of zoom you are using.
Show values on point	When enabled, you hover the mouse over the object icon to view the object name, type, and state.
Zoom the view	Click and drag the selection box in the map to enlarge the selected area.
Zoom in	Enlarges the map.
Zoom out	Decreases the size of the map.
Reset to initial resource	Returns the map to original object if you double-clicked on an icon to examine another object.
Resource detail	Changes the view in the main pane to the object details. You can use the Summary, Alerts, Analysis, and related tabs to troubleshoot the problem in more detail.

Table 5-12. Relationship Map (continued)

Option	Description
Show alerts	Opens a window that lists the alerts for the object you selected in the map.
Map	Topological view of the object and the related objects. Double-click on an object to see a relationship map for that object. The metric chart selector list is based on the object that is the focus of the map.

The chart options are used to limit the metric list.

Table 5-13. Metric Chart Selector

Option	Description
Show common metrics	Updates the list to show only the metrics that are available for the object type.
Show collecting metrics	Updates the list to display only the currently collected metrics for the object type.
Actions	Click the **Actions** icon to configure metric groups. Verify that you hold the PowerUser or administrator role. ■ **Add Group**. To add metrics or properties to the group, expand any of the metric groups, and drag one or more metrics to the group. ■ **Remove Group(s)**. To remove one or more groups. ■ **Rename Group**. To enter a new name for the group. ■ **Remove Metric(s) from Group(s)**. To remove one or more metrics or properties from one or more groups, hold down the Ctrl key, and select the metrics or properties that you want to remove.
Search	Use a word search to limit the number of items that appear in the list.
Metric list	Double-click a metric to populate the chart window. To populate the chart window with a separate chart for each of the metrics in the group, double-click a metric group.

To visualize the specific metric data over time, and compare the results for different metrics, select different combinations of options.

Table 5-14. Metric Chart Toolbar

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.
Y Axis	Shows or hides the Y-axis scale.

Table 5-14. Metric Chart Toolbar (continued)

Option	Description
Metric Chart	Shows or hides the line that connects the data points on the chart.
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.
Show 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 to Fit	Resets the chart to fit in the available space.
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.
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.
Show Data Values	Enables the data point tooltips if you switched to a zoom or pan option. **Show Data Point Tips** must be enabled.
Refresh Charts	Reloads the charts with current data.

Table 5-14. Metric Chart Toolbar (continued)

Option	Description
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.
Generate Dashboard	Saves the current charts as a dashboard.
Remove All	Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts.

Manage individual charts with the toolbar options.

Table 5-15. Individual Metric Charts Toolbar

Option	Description
Navigation	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.
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.
Move Down	Moves the chart down one position.
Move Up	Moves the chart up one position.
Close	Deletes the chart.

Host-Related Metrics

vRealize Operations Manager provides groups of metrics for selected hosts. Each group displays the most relevant metrics for the host to help you monitor your environment.

To display metric groups, select a host in the Environment Overview, and then select the **All Metrics** tab.

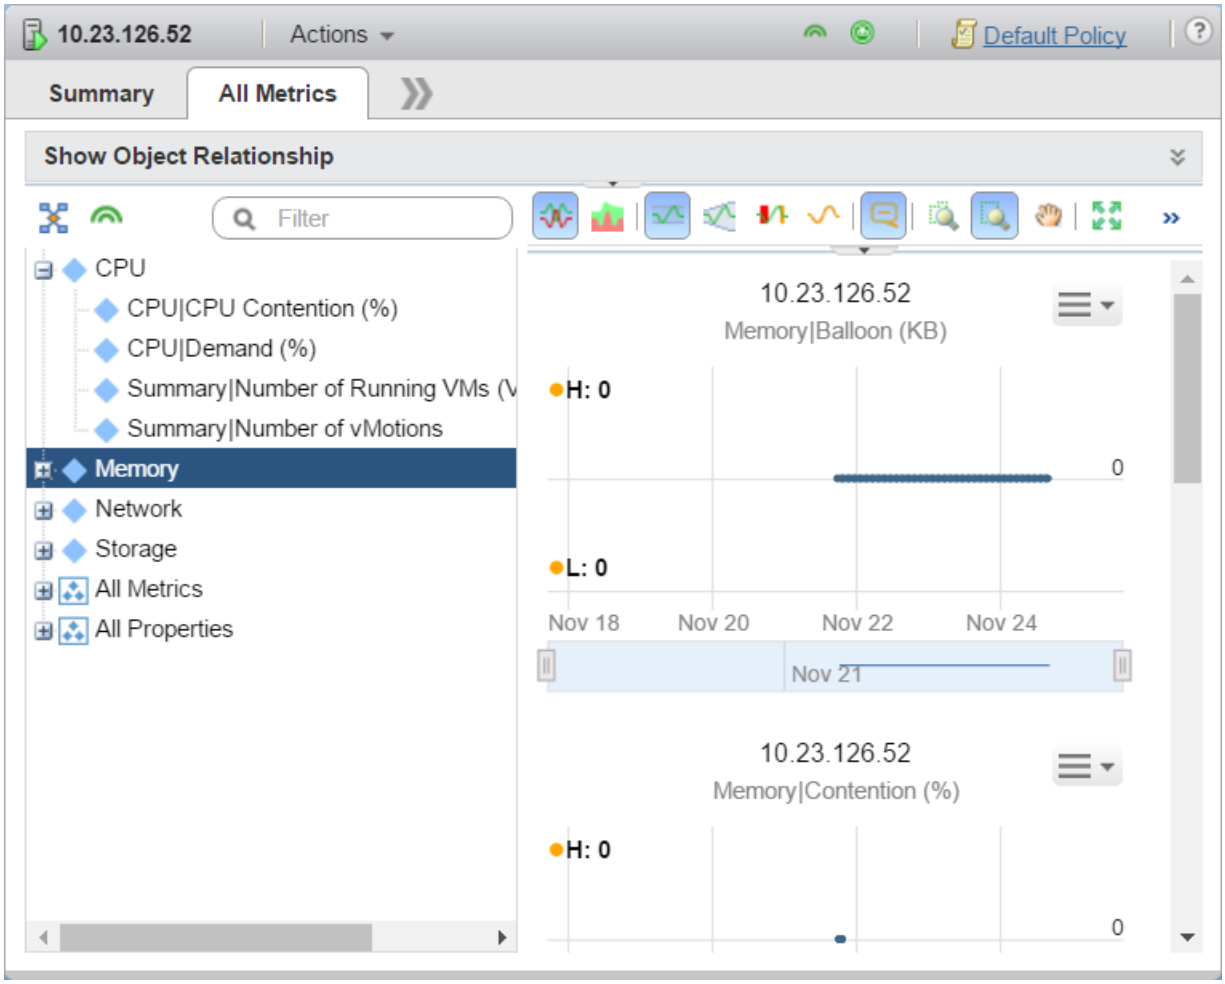

To display the metrics contained within a group, click the plus sign next to the group. You can double-click a group to populate the chart window with a separate chart for each of the metrics in the group. In the screenshot above, the metrics of the memory group populate the chart window.

Table 5-16. CPU Metric Group

Metric	Description
CPU|CPU contention (%)	This metric shows the percentage of time the VMs in the ESXi hosts are unable to run because they are contending for access to the physical CPUs. The number shown is the average number for all VMs. The number will be lower than the highest number experienced by the VM that is most impacted by CPU contention. Use this metric to verify if the host can serve all its VMs efficiently. Low contention means that the VM can access everything it demands to run smoothly. It means that the infrastructure is providing good service to the application team. When using this metric, ensure that the number is within your expectation. Look at both the relative number and the absolute number. Relative means a drastic change in value, meaning that the ESXi is unable to serve the VMs. Absolute means that the real value itself is high. Investigate why the number is high. One factor that impacts this metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz, the reduction in speed is accounted for because it shows that the VM is not running at full speed. This metric is calculated in the following way: **cpu|capacity_contention / (200 * summary|number_running_vcpus)**
CPU|Demand (%)	This metric shows the amount of CPU resources a VM would use if there were no CPU contention or CPU limit. This metric represents the average active CPU load for the past five minutes. Keep this number below 100% if you set the power management to maximum. This metric is calculated in the following way: **(cpu.demandmhz / cpu.capacity_provisioned)*100** .

Table 5-16. CPU Metric Group (continued)

Metric	Description
Summary|Number of running VMs	This metric shows the number of running VMs at a given point in time. The data is sampled every five minutes. A large number of running VMs might be a reason for CPU or memory spikes because more resources are used in the host. The number of running VMs gives you a good indicator of how many requests the ESXi host must juggle. Powered off VMs are not included because they do not impact ESXi performance. A change in the number of running VMs can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, because all the VMs will fail if the ESXi crashes. Use this metric to look for a correlation between spikes in the running VMs and spikes in other metrics such as CPU contention, or memory contention.
Summary|Number of vMotions	This metric shows the number of times a live migration (vMotion) with no VM downtime or service disruption took place in a host in the last (x) minutes. The number of vMotions is a good indicator of stability. In a healthy environment, this number is stable and relatively low. When using this metric, look for a correlation between vMotions and spikes in other metrics such as CPU contention and memory contention. Although the vMotion should not create any spikes, it is highly likely that some spikes in memory usage contention, and CPU demand and contention are experienced.

Table 5-17. Memory Metric Group

Metric	Description
Memory|Balloon (KB)	This metric shows the total amount of memory currently used by the VM memory control. Use this metric to monitor how much VM memory the ESXi has reclaimed through memory ballooning. The presence of ballooning indicates that the ESXi has been under memory pressure. ESXi activates ballooning when its consumed memory reaches a specific threshold. For example, in vRealize Operations Manager 6.0, the threshold is >98%. When using this metric, verify if the size of the ballooning is increasing. An increase in ballooning indicates that the lack of memory is not a one time occurrence, and that the memory shortage is worsening. Look for memory fluctuations which indicate that the VM required the ballooned out page. If the VM requests a ballooned out page, this translates into a memory performance problem for the VM because the page has to be returned from the disk. When the balloon target value is greater than the value shown by the metric, it means that there is more available memory that can be reclaimed.
Memory|Contention (%)	This metric shows the percentage of time VMs are waiting to access swapped memory. Use this metric to monitor ESXi memory swapping. A high value indicates that the ESXi is running low on memory, and a large amount of memory is being swapped.
Memory|Usage (%)	This metric shows the amount of physical memory actively used. The memory usage is displayed as a percentage of the total configured or available memory. This metric maps to the Consumed counter in vCenter. When the metric displays a high value, it indicates that the ESXi is using a large percentage of available memory. Check other memory-related metrics to see if the ESXi requires more memory.

Table 5-18. Network Metric Group

Metric	Description
Network I/O | Aggregate of all instances | Packet Dropped (%)	This metric shows the percentage of received and transmitted packets dropped in the collection interval. Use this metric to monitor the reliability and performance of the ESXi network. A high value indicates that the network is not reliable and performance decreases.
Network I/O | Aggregate of all instances | Packet Received per second	This metric shows the number of packets received in the collection interval. Use this metric to monitor the network usage of the ESXi.
Network I/O | Aggregate of all instances | Packet Transmitted per second	This metric shows the number of packets transmitted during the collection interval. Use this metric to monitor the network usage of the ESXi.

Table 5-19. Storage Metric Group

Metric	Description
Datastore I/O|Average observed virtual machine disk I/O workload	
Storage adapter|Aggregate of all instances|Read latency (ms)	This metric shows the average amount of time required for a read operation by all the storage adapters. Use this metric to monitor the read operation of the storage adapter. A high value indicates that the ESXi is experiencing storage read operation slowness. The total latency is the sum of kernel latency and device latency.
Storage adapter|Aggregate of all instances|Write latency (ms)	This metric shows the average amount of time required for a write operation by all the storage adapters. Use this metric to monitor the write operation performance of the storage adapter. A high value indicates that the ESXi is experiencing storage write operation slowness. The total latency is the sum of the kernel latency and device latency.

Analyzing the Resources in Your Environment

In addition to monitoring, vRealize Operations Manager provides you with powerful tools for analyzing the resources and the performance of your virtual environment.

You can use the Analysis tab to analyze the current condition of your virtual environment.

Analysis Badge Definitions

vRealize Operations Manager uses badges to visualize metrics to give you a high level view of the performance and the condition of your virtual environment.

The scores for the analysis badges are computed by the vCenter Server adapter, and others by the vRealize Operations Manager analytics algorithms.

Table 5-20. vRealize Operations Manager badges

Name	Icon	Description
Workload		The Workload badge combines metrics that show the demand for resources on an object as a single value. These metrics include CPU utilization, memory usage, and so on.
Anomalies		The Anomalies score is calculated using the total number of threshold violations for all metrics for the selected object and its child objects. A low Anomalies score indicates that an object is behaving according to its established historical parameters.
Faults		The Faults score is calculated based on events published by the vCenter Server. The scores are computed based on the severity of underlying problems. When more than one fault-related problem exists on the resource, the Faults score is based on the most severe problem.
Capacity		The Capacity badge represents the capability of your virtual environment to accommodate new virtual machines. vRealize Operations Manager calculates the Capacity score as a percentage of the remaining virtual machines count compared to the total number of virtual machines that can be deployed on the selected object.
Time Remaining		The Time Remaining score indicates how much time is remaining before the resources of the object exhaust. The Time Remaining score allows you to plan the provisioning of physical or virtual resources for the selected object, or reorganising the workload in your virtual environment.
Stress		The Stress score indicates the historic workload of the selected object. The Stress score is calculated as a ratio between the demand for resources and the usable capacity for a certain period.
Reclaimable Capacity		The Reclaimable Waste score indicates over-provisioning in your virtual infrastructure or for a specific object. It identifies the amount of resources that can be reclaimed and provisioned to other objects in your environment.

Table 5-20. vRealize Operations Manager badges (continued)

Name	Icon	Description
Density		The Density score indicates consolidation ratios, such as virtual machines per host, virtual CPUs per physical CPU, virtual memory per physical memory, and so on. You can use the Density score to achieve higher consolidation ratios and cost savings.
Compliance		The Compliance badge value is a score based on one or more compliance templates that you run in vRealize Operations Manager against the data collected from vRealize Operations Manager, datacenter, cluster, host system, virtual machine objects that are managed by vRealize Operations Manager and also by vRealize Configuration Manager if you have this adapter installed. The scores are calculated basedvRealize Configuration Manager on configured settings.

Badge Scores

The badge score ranges between 0 and 100. For Time Remaining, Capacity, Efficiency, and Density badges a score of 100 indicates, they are in good condition. For Workload, Anomalies, Faults, Stress, and Reclaimable Capacity badges, 100 indicates bad condition. The color is based on badge score thresholds that are set by thevRealize Operations Manager administrator. Each badge has default thresholds. See Policies for more information about configuring the thresholds for badge scores.

The badges do not indicate the power state of vSphere-related resources. For example, if a host is disconnected in vSphere, the Workload badge will show Unknown instead of Offline status.

Badge Color	Icon	Description
Green colored badge		The object is in normal state, based on the set thresholds. For example, by default, the green infrastructure Workload badge indicates a score above 76.
Yellow colored badge		The object is experiencing some level of problems. For example, by default the infrastructure yellow Workload badge indicates a score between 80 and 89.
Orange colored badge		The object might have serious problems or is approaching its capacity. For example, by default, the infrastructure orange Faults badge indicates a score between 50 and 74.
Red colored badge		The object is either not functioning properly or will stop functioning soon. Most of the metrics are beyond their thresholds. For example, by default, the infrastructure red Risk badge indicates a score 100.
Gray colored badge		No data is available for this object or the object is offline. For example, indicates that there is no data for the capacity remaining of the object.

The Workload Tab

Workload in vRealize Operations Manager is the demand for resources that an object requires versus the actual capacity that the object is able to access. The Workload badge value is a score based on how hard an object must work for resources. Use the Workload value as an investigative tool when you are researching capacity constraints or evaluating the general state of objects in your environment.

The Workload Badge

The vRealize Operations Manager Workload analysis badge indicates how hard an object must work for resources. vRealize Operations Manager indicates the workload by a colored icon that is based on the defined badge score thresholds.

The Workload score ranges from 0 (good) to over 100 (bad). The badge score thresholds can by modified by the vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Workload badge score.

Table 5-21. Object Workload States

Badge Icon	Description	User Action
	Workload on the object is not excessive.	No attention required.
	Object is experiencing some high-resource workloads.	Check the details and take appropriate action.
	Workload on the object is approaching its capacity in at least one area.	Check the details and take appropriate action as soon as possible.
	Workload on the object is at or over its capacity in one or more areas.	Act immediately to avoid or correct problems.
	The object is offline or no data is available.	

Where You Find the Workload Badge

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Workload** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-22. Workload Based on the Selected Inventory Object

Item	Description
Badge status	Workload status for the object based on the workload policy.
Workload Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Workload Breakdown	Breakdown of the current workload. The information displayed depends on the object type.
Workload in Related Objects	Workload status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.
Further Analysis	Lists further analysis options to troubleshoot workload issues related to the selected analysis badge. Further analysis uses the last 24 hours of data.
Object Resources	Configured resources for the object.
Workload Policy Settings	Policy settings shows what is going to be used for a workload.

The Anomalies Tab

The information shown in the **Anomalies** tab is the anomalies reported when metric values fall outside their normal range. The Anomalies score is the percentage of all metrics that have abnormal behavior. Use the Anomalies value as an investigative tool when you are researching abnormal behavior or evaluating the general state of objects in your environment.

vRealize Operations Manager calculates dynamic thresholds for each metric that is collected for an object. vRealize Operations Manager also analyzes the number of metrics that are violating their dynamic thresholds, to determine trends and normal levels of threshold violations. Based on these trends, the Anomalies analysis score is calculated using the total number of threshold violations for all metrics for the selected object and its child objects.

The Anomalies Badge

The vRealize Operations Manager Anomalies badge score represents how abnormal the behavior of the object is, based on its historical metrics data. vRealize Operations Manager indicates Anomalies using a colored icon that is based on the defined badge score thresholds.

When evaluating badge scores, a high number of anomalies might indicate a potential issue. A low Anomalies score indicates that an object is behaving in accordance with its established historical parameters. Most or all of the object's metrics, especially its KPIs, are within their thresholds. Because changes in behavior often indicate developing problems, if the metrics of an object go outside the calculated thresholds, the anomalies score for the object increases. As more metrics breach the thresholds, anomalies continue to increase.

Violations by KPI metrics increase the Anomalies score more than violations by non-KPI metrics. A high number of anomalies usually indicates a problem, or at least a situation that requires your attention.

Anomalies involve the number of statistics that fall outside of the expected behavior trends, whereas Workload involves an absolute measurement of how hard an object works for resources. Both Anomalies and Workload are useful when you are attempting to find a probable cause and when troubleshooting performance problems.

The Anomalies score ranges between 0 (good) and 100 (bad). The badge score thresholds can by modified by an vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Anomalies badge score.

Table 5-23. Object Anomalies States

Badge Icon	Description	User Action
	The Anomalies score is normal.	No attention required.
	The Anomalies score exceeds the normal range.	Check the details and take appropriate action.
	The Anomalies score is very high.	Check the details and take appropriate action as soon as possible.
	Most of the metrics are beyond their thresholds. This object might not be working properly or might stop working soon.	Act immediately to avoid or correct problems.
	The object is offline or no data is available.	

Where You Find the Anomalies Badge

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Anomalies** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-24. Anomalies Based on the Selected Inventory Object

Item	Description
Badge status	Anomalies status for the object based on the anomalies policy.
Anomalies Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Anomalies Breakdown	Breakdown of the compliance standards by alert. To see the violated standards, click the row in the standard table. To limit the standards list, use the following buttons. ■ Violated Standards. Displays only the standards alerts where at least on symptom is triggered. ■ All Standards. Displays all the standards alerts.
Anomalies in Related Objects	Anomalies of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.
Object Resources	Configured resources for the object.

The Faults Tab

The information shown in the **Faults** tab is a combination of the availability of the selected object and any configuration issues related to it. Each fault has a severity level. The Faults score is the severity level of the worst open fault, and is calculated based on events published by vCenter Server. The higher the Faults score, the lower the resulting health for that object.

The Faults Badge

The Faults score calculation includes events such as loss of redundancy in NICs or HBAs, memory checksum errors, HA failover problems, Common Information Model (CIM) events, and so on. Faults are included in the health score because they require immediate resolution, whereas items that contribute to the risk score might not be immediate, although requiring your attention.

Each object in vRealize Operations Manager has a faults score ranging from 0 (no faults) to 100 (critical faults). The scores are computed based on the severity of the underlying problems. When more than one fault-related problem exists on a resource, the faults score is based on the most severe problem.

Unlike other badges in vRealize Operations Manager, no alert is generated from the threshold score of the faults badge. Instead, each problem generates its own fault alert, and resolution of the problem both clears or cancels the alert and lowers the badge score.

The Faults score ranges between 0 (good) and 100 (bad). The badge score thresholds can by modified by an vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Anomalies badge score.

Table 5-25. Object Faults States

Badge Icon	Description	User Action
	No faults are registered on the selected object.	No attention required.
	Faults of low importance are registered on the selected object.	Check the details and take appropriate action.
	Faults of high importance are registered on the selected object.	Check the details and take appropriate action as soon as possible.
	Faults of critical importance are registered on the selected object.	Act immediately to avoid or correct problems.
	The object is offline or no data is available.	

Where You Find the Faults Badge

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Faults** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-26. Faults Based on the Selected Inventory Object

Item	Description
Badge status	Faults status for the object based on the combination of the availability of an object and any configuration issues.
Faults Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Faults Breakdown	Breakdown of the current faults affecting the group objects.
Faults in Related Objects	Faults status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

The Capacity Remaining Tab

The **Capacity Remaining** tab indicates the unused capacity of your virtual environment to accommodate new virtual machines. vRealize Operations Manager calculates the Capacity

Remaining score as a percentage of the remaining capacity count, compared to the total amount of capacity that can be deployed on the selected object.

The score is based on the current amount of unused resources and the average virtual machine profile for the last *n* weeks. The remaining virtual machines count is a function of the same compute resources of CPU, memory, disk I/O, net I/O, and disk space that are used to calculate the Time Remaining score.

Note Small, medium, average, and large virtual machine profiles cannot be computed for objects that do not have active child virtual machines. If the child virtual machines are powered off, you might see a '?' instead of a numerical value.

The Capacity Remaining Badge

The Capacity Remaining score ranges between 0 (bad) and 100 (good).

The badge score thresholds can by modified by a vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Capacity remaining badge score.

Table 5-27. Object Capacity States

Icon	Description	User Action
	The capacity remaining for the object is at normal level.	No attention required.
	The capacity remaining for the object is less than the normal level.	Check the details and take appropriate action.
	The capacity remaining for the object is at seriously low level.	Check the details and take appropriate action as soon as possible.
	The object is expected to run out of capacity soon or has already run out of capacity.	Act immediately to avoid or correct problems.
	The object is offline or no data is available for any of the metrics for the time period.	

Where You Find the Capacity Remaining Badge

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Capacity Remaining** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-28. Capacity Remaining Based on the Selected Inventory Object

Item	Description
Badge status	Capacity remaining status for the object.
Capacity Remaining Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Capacity Remaining Breakdown	Breakdown of the capacity remaining for objects. The data range considered for computing the capacity remaining for the resource containers is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Capacity Remaining in Related Objects	Capacity remaining status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

Time Remaining Tab

The **Time Remaining** tab indicates how much time is remaining before the resources of a selected object are exhausted. The Time Remaining score is the number of days until the maximum capacity is reached, minus the provisioning time buffer, based on your current consumption trend. The Time Remaining score allows you to plan the provisioning of physical or virtual resources for the selected object, or change the workload to adjust the needs of the resources in your virtual environment.

The Time Remaining score is calculated per resource type for an object. For example, CPU usage or disk I/O is based on the historical data for the object type. vRealize Operations Manager calculates the Time Remaining score as a percentage of time that is remaining for each compute resource compared to the provisioning buffer you set in the Configuration dialog box. By default, the Time Remaining score provisioning buffer is 30 days. If even one of the compute resources has less capacity than the provisioned buffer, the Time Remaining score is 0.

For example, if the provisioning buffer is set to 30 days, and the object that you selected has CPU resources for 81 days, memory resources for 5 days, disk I/O resources for 200 days, and network I/O resources for more than one year, the Time Remaining score is 0, because one of the resources has capacity for less than 30 days.

The Time Remaining Badge

The Time Remaining score ranges between 0 (bad) and 100 (good). The badge score thresholds can by modified by the vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Time Remaining badge score.

Table 5-29. Time Remaining States

Badge Icon	Description	User Action
	The number of days that remain is much higher than the score provisioning buffer you specified.	No attention required.
	The number of days that remain is higher than the score provisioning buffer, but is less than two times the buffer you specified.	Check the details and take appropriate action.
	The number of days that remain is higher than the score provisioning buffer, but approaches the buffer you specified.	Check the details and take appropriate action as soon as possible.
	The number of days that remain is lower than the score provisioning buffer you specified. The selected object might have exhausted some of its resources or will exhaust them soon.	Act immediately to avoid or correct problems.
	The object is offline or no data is available for the Time Remaining score.	

Where You Find the Time Remaining Score

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Time Remaining** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-30. Time Remaining Based on the Selected Inventory Object

Item	Description
Badge status	Time remaining status for the object.
Time Remaining Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.

Table 5-30. Time Remaining Based on the Selected Inventory Object (continued)

Item	Description
Time Remaining Breakdown	Breakdown of the time remaining for objects.
Time Remaining in Related Objects	Time remaining status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

The Stress Tab

The **Stress** tab indicates represents how vRealize Operations Manager calculates how much an object demands over a period of time. This analysis compares an object's workload against its capacity. The Stress score helps you identify hosts and virtual machines that do not have enough resources allocated, or hosts that are running too many virtual machines.

Stress accumulates when workload exceeds the specified stress line. The Stress score is the percentage of the stress zone area with stress in your selected time sample. A high Stress score does not imply a current performance problem, but highlights potential for future performance problems.

The Stress Badge

The Stress score ranges between 0 (good) and 100 (bad). The badge score thresholds can by modified by the vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Stress badge score.

Table 5-31. Stress States

Badge Icon	Description	User Action
	The Stress score is normal.	No attention required.
	Some of the object resources are not enough to meet the demands.	Check the details and take appropriate action.
	The object is experiencing regular resource shortage.	Check the details and take appropriate action as soon as possible.
	Most of the resources on the object are constantly insufficient. The object might stop functioning properly.	Act immediately to avoid or correct problems.
	The object is offline or no data is available for the Stress score.	

Where You Find the Stress Score

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Stress** button.

- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-32. Stress Based on the Selected Inventory Object

Item	Description
Badge status	Stress status for the object.
Stress Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Stress Breakdown	Breakdown of stress. The data range considered for computing the stress for the resource containers is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Stress in Related Objects	Stress status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

The Reclaimable Capacity Tab

The **Reclaimable Capacity** tab indicates the amount of provisioned capacity that you can reclaim and provision to other objects in your environment without causing stress or performance degradation. Reclaimable Capacity is calculated for each resource type, such as CPU, memory, and disk, for each object in the environment.

For groups, Reclaimable Capacity is the amount of disk space that can be reclaimed from the virtual machines in the group that are considered waste, based on the policy settings for the powered off and idle state. If the virtual machine is idle, all its resources are considered reclaimable. If a group does not contain any virtual machines, but contains datastores, the value for Reclaimable Capacity is 0, even if the datastore contains virtual machines that are wasting resources based on the **Powered off and idle VMs** settings.

For more information about Reclaimable Capacity as it pertains to policy settings, see Policy Reclaimable Capacity Element.

The Reclaimable Capacity Badge

The Reclaimable Capacity analysis badge value is a score based on the percentage of your total capacity that could be re-purposed.

The Reclaimable Capacity score ranges between 0 (good) and 100 (bad). The badge score thresholds can by modified by the vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Reclaimable Capacity badge score.

Table 5-33. Reclaimable Waste States

Badge Icon	Description	User Action
	No resources are wasted on the selected object.	No attention required.
	Some resource can be used better.	Check the details and take appropriate action.
	Many resources are underused.	Check the details and take appropriate action as soon as possible.
	Most of the resources on the selected object are wasted.	Act immediately to avoid or correct problems.
	The object is offline or no data is available for any of the metrics for the time period.	

Where You find the Reclaimable Capacity Tab

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Reclaimable Capacity** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-34. Reclaimable Capacity Based on the Selected Inventory Object

Item	Description
Badge status	Reclaimable Capacity status for the object.
Reclaimable Capacity Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Reclaimable Capacity Breakdown	Breakdown of reclaimable capacity by object. The data range considered for computing the reclaimable capacity for the resource containers is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Reclaimable Capacity in Related Objects	Reclaimable Capacity status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

The Density Tab

The **Density** tab indicates the consolidation ratios, such as virtual machines per host, virtual CPUs per physical CPU, virtual memory per physical memory, and so on. You can use the Density score to achieve higher consolidation ratios and cost savings.

When you understand the behavior and performance of your virtual machines and applications, you can maximize the consolidation in your virtual environment without affecting performance or service-level agreements. Density analytics determine an optimal child to parent consolidation ratio.

The Density Badge

The Density badge value is a score based on the percentage alignment of your actual consolidation ratio to optimal.

The Density score ranges between 0 (bad) and 100 (good). The badge score thresholds can by modified by the vRealize Operations Manager administrator.

See Policies to configure symptom thresholds for the Density badge score.

Table 5-35. Object Density States

Badge Icon	Description	User Action
	The resource consolidation is good.	No attention required.
	Some resources are not fully consolidated.	Check the details and take appropriate action.
	The consolidation for many resources is low.	Check the details and take appropriate action as soon as possible.
	The resource consolidation is extremely low.	Act immediately to avoid or correct problems.
	The object is offline or no data is available for any of the metrics for the time period.	

Where You Find Density Information

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Density** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-36. Density Based on the Selected Inventory Object

Item	Description
Badge status	Density status for the object based on the most critical of the violated standards.
Density Trend	How the badge value for the object has trended over time. This trend view allows you to see behavior over time and to identify when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.

Table 5-36. Density Based on the Selected Inventory Object (continued)

Item	Description
Density Breakdown	Breakdown of the density percentage based on the percentage of the actual consolidation ratio to the optimal ratio. The data range considered for computing the density for the resource containers is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Density in Related Objects	Density status of the related objects. Use the related objects to determine if any problems are affecting only the current object or are related objects experiencing problems.

Compliance Tab

The **Compliance** tab provides analysis based on the vRealize Operations Manager alerts that are configured with the alert subtype of Compliance. You use the compliance value as an investigative tool when you evaluate the state of objects in your environment, or when you research the root cause of a problem.

You can use the alert-based compliance that vRealize Operations Manager provides to ensure compliance of your vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches. If you also use vRealize Configuration Manager in your environment, you can add the vRealize Configuration Manager adapter to vRealize Operations Manager. The vRealize Configuration Manager adapter provides vRealize Configuration Manager compliance information in place of the alert-based compliance.

The compliance alerts, which have the subtype named Compliance, include one or more symptoms that represent the compliance rules. Compliance alerts that trigger appear on the **Compliance** tab as a 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 triggers for any of the alerts in the standard, the triggered rule violates the standard and affects the badge score that appears on the **Compliance** tab.

The Compliance Badge

To calculate the compliance badge score, vRealize Operations Manager uses the compliance percentage, total count of symptoms, and the count of triggered symptoms.

The compliance percentage calculation is:

```
100 – ((triggered symptom count(TR)/total symptom count(TS)) * 100)
```

In this calculation, the following statements are true.

- The total symptom count includes all symptoms in all active compliance alerts.
- The triggered symptom count includes all triggered symptoms in all active compliance alerts.

The threshold values determine the following compliance scores:

- 100 indicates a good score

- 51-99 indicates a warning score
- 26-50 indicates an immediate score
- 0-25 indicates a critical score

To enable alert-based compliance, you must customize a policy. If the compliance alerts are not enabled, the Compliance badge value is 100 and is green, and no violations exist in the list of violated standards. For example, the VMware vSphere solutions provide the alerts for the ESXi host and virtual machine sections of the *vSphere Hardening Guide*.

To customize policies to enable alert-based compliance, see Customize a Policy to Enable the vSphere Hardening Guide Alerts.

Where You Find Compliance Based on vRealize Operations Manager Alerts

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Analysis** tab, then select the **Compliance** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-37. Compliance Based on vRealize Operations Manager Alerts Options

Item	Description
Badge status and score	Compliance status and score for the object based on the most critical of the violated standards. The badge displays one of the following values: ■ 100 indicates a good score, with no triggered symptoms in the compliance alerts. The badge color is green. ■ 51 to 99 indicates a warning that some symptoms triggered in the compliance alerts. The badge color is yellow. ■ 26 to 50 indicates an immediate score, because numerous symptoms triggered in the compliance alerts. You must take action immediately. The badge color is orange. ■ 0 to 25 indicates a critical score. You must take action immediately. The badge color is red.
Compliance Trend	Indicates how the badge value for the object has changed over time. The trend displays the behavior over time, and identifies when a change in a badge value indicates a change on the object. The trend data time value is based on the Data Range setting, which is defined in the Time analysis settings for the policy that is associated with the object.
Compliance Breakdown	Displays the breakdown of the compliance standards by alert. To see the violated standards, click the row in the table of standards. To focus your view on the standards list, click the following buttons. ■ Violated Standards. Displays only the alerts in the standards where at least one symptom is triggered. ■ All Standards. Displays all alerts in the standards.

Table 5-37. Compliance Based on vRealize Operations Manager Alerts Options (continued)

Item	Description
Violated rules list	Violated rules are the symptoms defined in the compliance alert. If you click the standard, the rules for the standard appear. If a symptom triggered, the rule is considered to be violated. To focus the rules list, use the following buttons. ■ Violated Rules. Displays only the triggered symptoms. ■ All Rules. Displays triggered and untriggered symptoms.
Compliance in Related Objects	Displays the compliance status of the related objects. Use the related objects to determine whether any problems are only affecting the current object, or if related objects are experiencing problems.
Object Resources	Displays the configured resources for the object.

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

Customize a Policy to Enable the *vSphere Hardening Guide* Alerts

The *VMware vSphere Hardening Guide* alerts notify you when settings or properties on your vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches are not configured in compliance with the guide. To have vRealize Operations Manager assess your objects against the compliance alerts, you must override the policy state so that the setting named Local is enabled for each alert.

The alert-based compliance works after you enable the *VMware vSphere Hardening Guide* alerts. The VMware vSphere Hardening Guide checks the collected data to determine whether the settings are configured correctly so that your objects operate in a secure manner.

Prerequisites

Verify that your instance of vRealize Operations Manager includes the Default Policy and one or more other policies. See Default Policy in vRealize Operations Manager.

Procedure

1. In the menu, click **Administration**.
2. Click **Policies** and click the **Policy Library** tab.
3. Expand **Base Settings**, click the policy to customize it, and click the pencil to edit the selected policy.
4. In the Edit Monitoring Policy workspace, click **Alert / Symptom Definitions**.
5. Select Alert Definitions pane to display and examine the compliance alerts and enter **`hardening`** in the text box.

   Table 5-38. Compliance Alerts

Compliance Alerts	Support for *vSphere Hardening Guide* Version
ESXi host is violating *vSphere Hardening Guide*	5.5 and 6.0
vCenter Server is violating *vSphere Hardening Guide*	6.0

Table 5-38. Compliance Alerts (continued)

Compliance Alerts	Support for *vSphere Hardening Guide* Version
Virtual machine is violating Risk Profile 1 in *vSphere Hardening Guide*	5.5 and 6.0
Virtual machine is violating Risk Profile 2 in *vSphere Hardening Guide*	5.5 and 6.0
Virtual machine is violating Risk Profile 3 in *vSphere Hardening Guide*	5.5 and 6.0
vSphere Distributed Port Group is violating *vSphere Hardening Guide*	6.0
vSphere Distributed Virtual Switch is violating *vSphere Hardening Guide*	6.0

6 For each compliance alert, click the **State** drop-down menu and click **Local**.

7 To save your updates to the policy, click **Save**.

Results

You have enabled the alerts and the associated symptom definitions. When the configured policy is applied to objects, it becomes active. When the configured symptom definitions become true for your vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches, vRealize Operations Manager generates the *VMware vSphere Hardening Guide* alerts.

What to do next

Review the **Compliance** tab to determine whether your objects are in compliance. For an example, see User Scenario: Ensure Host Objects Comply With Alert-Based Compliance Rules.

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

User Scenario: Ensure Host Objects Comply With Alert-Based Compliance Rules

As a virtual infrastructure administrator, you use vRealize Operations Manager to monitor the objects in your environment, including vCenter Server instances and ESXi hosts, on which run your virtual machines. You review the **Compliance** tab for your hosts and discover that one of your hosts is violating the *VMware vSphere Hardening Guide* standard. You must identify and fix the problems.

vRealize Operations Manager includes alert-based compliance from the *VMware vSphere Hardening Guide*.

In this scenario, you resolve a violated rule on your host, and another violated rule on one of your virtual machines. In your own scenario, you would repeat this procedure for any other violated rules.

vRealize Operations Manager assesses vSphere 6.0 objects against 6.0 rules, and vSphere 5.5 objects against 5.5 rules.

Prerequisites

- Verify that you can open an XLSX file on the machine that you are using to access vRealize Operations Manager.
- Enable the *vSphere Hardening Guide* alerts so that the alert-based compliance is active in your environment. See Define Monitoring Goals for vRealize Operations Manager Solutions.

Procedure

1. In the menu, click **Environment**.
2. Browse to a host object.

   If you had created an object group to manage your hosts, you would select a host in the group.
3. With the host as the focus, click the **Analysis** tab and click the **Compliance** tab.

   The Compliance badge displays a value other than 100 or green.
4. Click the violated standard named **ESXi Host is violating vSphere Hardening Guide**.

   The Compliance Breakdown area expands to display all the violated rules, including violations for vSphere 6.0 objects and 5.5 objects.
5. Review the page to determine the criticality and pervasiveness of the noncompliant standards for this host and your environment.

Option	Evaluation
Compliance Breakdown	What is the number and criticality of the violated rules for the host? How many of the violated rules are critical and must be addressed?
Compliance in Related Objects	Are other hosts in a similar compliance state? Are any child objects out of compliance?
Host System Resources	Is the host configured as you expect?

   The page indicates that you must resolve the violated rule named `ESXi Host is violating vSphere Hardening Guide`.
6. Click the **Alerts** tab.

   The compliance standards are based on alerts, which can include recommendations. For example, the alert named `ESXi Host is violating vSphere Hardening Guide` includes a recommendation that links to the *VMware vSphere Hardening Guide*.
7. On the **Alerts** tab, click the alert named **ESXi Host is violating vSphere Hardening Guide**.

   The **Alert Details Summary** tab displays the violated rules as symptoms, and includes the recommendations to resolve the alert.

8 In the Recommendations area, click the link to the *vSphere Hardening Guides* at: http://www.vmware.com/security/hardening-guides.html, and click the link to the version you need.

   The *vSphere Hardening Guide* downloads as an Excel spreadsheet to the machine you are using to access vRealize Operations Manager.

9 You see that vRealize Operations Manager identified that one of the virtual machines is violating a DCUI rule, so you locate the compliance rule and the remediation method.

   - For vSphere 6.0 objects, in the 6.0 version of the *vSphere Hardening Guide*, locate the rule named `Set DCUI.Access to allow trusted users to override lockdown mode`.
   - For vSphere 5.5 objects, in the 5.5 version of the *vSphere Hardening Guide*, click the **ESXi** tab and locate the rule named `Disable DCUI to prevent local administrative control`.

10 Review information about the rule in the *vSphere Hardening Guide*, and implement the remediation method.

Results

You identified and resolved violated compliance rules that triggered on your host and virtual machine. After you remediate the violated rules, as described in the *vSphere Hardening Guide*, wait for vRealize Operations Manager to run several collection cycles. After several collection cycles, the violated rules no longer appear in the list of violated standards.

Using Troubleshooting Tools to Resolve Problems

The data provided in the **Alerts**, **Symptoms**, **Timeline**, **Events**, and **All Metrics** tabs help you identify the root cause of a complex problem.

You can use the troubleshooting tabs individually or as part of a workflow to resolve problems. Each of the tabs displays the collected data in a different way. Sometimes, as you are troubleshooting problems, you move directly from an analysis tab to the **All Metrics** tab. Under other circumstances, the **Timeline** tab might provide the information that you need.

Symptoms Tab Overview

You can view a list of triggered symptoms for the selected object. You use the symptoms when you are troubleshooting problems with an object.

The **Symptoms** tab displays all the triggered symptoms for the currently selected object. A review of the triggered symptoms provides you with a list of the problems that the currently selected object is experiencing. If you need to better understand which symptoms are associated with currently generated alerts, go to the **Alerts** tab for the object.

As you evaluate the triggered symptoms, consider the time at which they were created and the configuration information and trend charts, where applicable.

Symptoms Tab

The symptoms tab includes all the symptoms triggered for the current object. Use the symptom list to identify problems with an object so that you can resolve alerts generated for the object.

How the Symptoms Work

The list is the active triggered symptoms for an object, either as part of a generated alert or as a triggered symptom that is not included in an alert. This complete symptom list is useful for identifying problems that occur on an object but are not currently included in your alert definitions.

Click a symptom in the list to display the symptom details. An arrow in each column heading enables you to order the list in ascending or descending order. You can select multiple rows in the list using Shift+click, Control+click.

Where You Find the Symptoms Tab

- In the menu, select **Environment**, then select a group, application, or inventory object. Click the **object** to bring up the object's **Summary** tab. Click the **Alerts > Symptoms** tabs.
- In the menu, select **Search** and locate the object of interest. Click the **object** to bring up the object's **Summary** tab. Click the **Alerts > Symptoms** tabs.

Table 5-39. Symptoms Data Grid

Option	Description
Criticality	Criticality is the level of importance of a symptom in your environment. The level is based on the level assigned when the symptom was created. The possible values include: ■ Critical ■ Immediate ■ Warning ■ Information
Symptom	Name of the triggered symptom.
Status	Current state of the symptom. Possible values are Active or Inactive
Created On	Date and time when the alert was generated.
Canceled On	Date and time when the symptom was canceled.
Information	Information about the triggering condition for the symptom, including the trend and current value. The sparkline displays a range of data that includes six hours before the symptom update time and one hour after the update time.

Table 5-40. Filters

Filtering options	Limits the list of symptoms to those matching the filter you select. Some filters are similar to data grid headings: Symptom, Status, Criticality, Created on, Canceled on.
Triggered On	Name of the object for which the symptom was generated. Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object.

Timeline Tab Overview

The timeline provides a view of the triggered symptoms, generated alerts, and events for an object over a period of time. You use the timeline to identify common trends over time that are contributing to the current status of objects in your environment.

The timeline provides a three-tier scrolling mechanism that you can use to move quickly through large spans of time, or slowly and minutely through individual hours when you are focusing on a particular period of time. To ensure that you have the data that you need, configure the Date Controls to encompass the problem you are investigating.

It is not always effective to investigate a problem on an individual object by looking only at the object. Use the parent, children, and peer options to examine the object in a broader environmental context. This context often reveals unexpected influences or consequences for the problem.

The timeline is a tool that provides you a graphical view of patterns. If a symptom is triggered and canceled by the system at various intervals over time, you can compare the event to other changes to the object or to the related objects. These changes might be the root cause of the problem.

Events Timeline Tab

The generated alerts, triggered symptoms, and change events for the current object over time appear on the **Timeline** tab. You use the timeline to identify common trends over time that are contributing to the current status of objects in your environment.

How the Events Timeline Works

The timeline view includes alerts, symptoms, and events for the selected object for the last 6 hours. To view the data for a particular time, click on the timeline in one of the 3 tiers and move your mouse to the left to see data from the past or to the right to move back to the present.

The view is limited to approximately 50 alerts, symptoms, and events. If your timeline includes more than this number, you can use the toolbar options to remove data from the timeline until it contains data that you find useful for your investigation.

Where You Find the Events Timeline

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **object** to bring up the object's **Summary** tab. Click the **Events > Timeline** tabs.

- In the menu, select **Search** and locate the object of interest. Click the **object** to bring up the object's **Summary** tab. Click the **Events > Timeline** tabs.

Table 5-41. View From Menu

Option	Description
Self	Shows or hides events for the current object.
Peer	Shows or hides events for objects like the impacted object.
Parents <options>	Shows or hides events for the parent, grandparent, and so on, objects of the current object.
Children <options>	Shows or hides the events for the descendants of the impacted object.

Table 5-42. Alert Filters

Option	Description
Criticality <options>	Limits the alerts to those matching the selected criticality level. If no criticality is selected, all alerts are displayed.
Status <options>	Limits the alerts in the chart to the canceled or active alerts. If no status is selected, all alerts are displayed. This option applies only to alerts, not to fault and change events. Change events and active faults are always displayed in the chart.
Alert Type <options>	Select one or more alert types. The types are assigned when the alert is defined. If no type is selected, all alerts are displayed.

Table 5-43. Event Filters

Option	Description
Dynamic Threshold Violation	vRealize Operations Manager calculates dynamic thresholds for each metric that is collected for an object based on policies set.
Hard Threshold Violation	Events that represent a hard threshold violation, based on policies set. The system analyses the number of metrics that are violating their hard thresholds to determine trends.
Data Availability	Events reflecting data store performance. Data availability is the capacity to provide data on demand to users and applications.
System Degradation	Events that reflect negative impacts on system performance.
Environment	Events indicating a change in the environment.
Change	Shows or hides the change events. Change events are changes to the object that might or might not result in an alert.

Table 5-43. Event Filters (continued)

Option	Description
Notification	Routine notification events.
Fault	Events indicating any observed behaviour that differs from the expected one.

Table 5-44. Date Controls, Data Values, Events Chart

Option	Description
Date Controls	Limits the data in the chart to the selected time frame.
Data Values	When you click a data point, the event is highlighted in the event data grid.
Events chart	Shows the events and alerts over time by criticality, and other data options you select in the toolbar.

Events Tab Overview

Events are changes in vRealize Operations Manager metrics that reflect changes that occurred on managed objects because of user actions, system actions, triggered symptoms, or generated alerts on an object. You use the **Events** tab to compare the occurrence of events with the generated alerts to determine if a change on your managed object contributed to the root cause of the alert or other problems with the object.

Events can occur on any object, not just the one listed.

The following vCenter Server activities are some of the activities that generate vRealize Operations Manager events:

- Powering a virtual machine on or off
- Creating a virtual machine
- Installing VMware Tools on the guest OS of a virtual machine
- Adding a newly configured ESX/ESXi system to a vCenter Server system

Depending on alert definitions, these events might generate alerts.

If you monitor the same virtual machines with other applications that provide information to vRealize Operations Manager, and the adapters for those applications are configured to provide change events, the **Events** tab includes certain change events that occur on the monitored objects. These change events might provide further insight into the cause of problems that you are investigating.

Events Tab

An event is any change to an object that is identified by a change in the vRealize Operations Manager metrics for that object. You can compare changes to an object with symptoms and other data to identify a possible cause for a generated alert.

How the Events Tab Works

If you arrive at the Events tab from the Alerts page or tab, the Events tab opens with the timeline centered on the moment the alert occurred for the selected object.

You can configure the chart to display various combinations of data, allowing you to identify events that contribute to the alert you are investigating. Use the range selectors to shift the larger time frame in the timeline, then click and drag on the graph area to zoom in on a specific period. Click the data points on the graph to see pop-up descriptions of the various events.

Click the **Actions** menu to open an external application, for example, vSphere Client.

Where You Find the Events Tab

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **object** to bring up the object's **Summary** tab. Click the **Events > Events** tabs.
- In the menu, select **Search** and locate the object of interest. Click the **object** to bring up the object's **Summary** tab. Click the **Events > Events** tabs.
- In the menu, click **Alerts**, then click an **alert** of interest to display the alert details on the right. Click **View events**. The object that triggered the alert is displayed with associated events.

Table 5-45. View From

Option	Description
Self	Shows or hides events for the current object.
Peer	Shows or hides events for objects like the impacted object.
Parents <options>	Shows or hides events for the parent, grandparent, and so on, objects of the current object.
Children <options>	Shows or hides the events for the descendants of the impacted object.

Table 5-46. Alert Filters

Option	Description
Criticality <options>	Limits the alerts to those matching the selected criticality level. If no criticality is selected, all alerts are displayed.
Status <options>	Limits the alerts in the chart to the canceled or active alerts. If no status is selected, all alerts are displayed. This option applies only to alerts, not to fault and change events. Change events and active faults are always displayed in the chart.
Alert Type <options>	Select one or more alert types. The types are assigned when the alert is defined. If no type is selected, all alerts are displayed.

Table 5-47. Event Filters

Option	Description
Dynamic Threshold Violation	vRealize Operations Manager calculates dynamic thresholds for each metric that is collected for an object based on policies set.
Hard Threshold Violation	Events that represent a hard threshold violation, based on policies set. The system analyses the number of metrics that are violating their hard thresholds to determine trends.
Data Availability	Events reflecting data store performance. Data availability is the capacity to provide data on demand to users and applications.
System Degradation	Events that reflect negative impacts on system performance.
Environment	Events indicating a change in the environment.
Change	Shows or hides the change events. Change events are changes to the object that might or might not result in an alert.
Notification	Routine notification events.
Fault	Events indicating any observed behaviour that differs from the expected one.

Table 5-48. Date Controls, Events Chart, Events Data Grid

Option	Description
Date Controls	Limits the data in the chart to the selected time frame.
Events chart	Shows the events and alerts over time by criticality, and other data options you select in the toolbar.
Events data grid	Shows a list of events when you select at least one of the following display options: ■ Self ■ Parent ■ Child ■ Peer

Creating and Using Object Details

The views and heat map details provide you with specific data about the object. You use this information to evaluate problems in more detail. If the current views or heat maps do not provide the information that you need, you can create one to use as tool as you investigate your specific problem.

Details Views Tab

The **Views** tab is divided into two panels. The bottom panel updates, depending on what you select on the top panel.

In the top panel you can create, edit, delete, clone, export, and import views. The views list depend on the object you select from the environment. Each view is associated with an object. For example, the predefined VM inventory - Memory list view is available when you select a host.

You can limit the views list by adding a filter from the right side of the panel. Each of the provided filter groups limits the list by the word you type. For example, if you select **Description** and type **`my view`**, the listed views are all views that are applicable for the selected object and contain *my view* in the description.

Table 5-49. Views List Table Columns

Column	Description
Name	Name of the view.
Type	Type of the view. A view type is the way the collected information for the object is presented.
Description	Description of the view as it is defined when the view is created.
Subject	Object type with which a view is associated.
Owner	Owner of the view is the user, who created it or edited it for the last time.

In the bottom panel of the **Views** tab you can see the data of the object, calculated by a selected view from the top panel. For example, if the selected object is a host and you select Virtual Machine Configuration Summary List View, the result is a list of all the virtual machines on that host, and their data calculated by the view.

For Trend views, you can select a parent object and see the data of the associated child objects and metrics in the bottom panel of the **Views** tab.

Where You Find the Details View Tab

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Details** tab, then select the **Views** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Working with Heat Maps

With the vRealize Operations Manager heat map feature, you can locate trouble areas based on the metric values for objects in your virtual infrastructure. vRealize Operations Manager uses analytics algorithms that you can use to compare the performance of objects across the virtual infrastructure in real time using heat maps.

You can use predefined heat maps or create your own custom heat maps to compare the metric values of objects in your virtual environment. vRealize Operations Manager has predefined heat maps on the **Details** tab that you can use to compare commonly used metrics. You can use this data to plan to reduce waste and increase capacity in the virtual infrastructure.

What a Heat Map Shows

A heat map contains rectangles of different sizes and colors, and each rectangle represents an object in your virtual environment. The color of the rectangle represents the value of one metric, and the size of the rectangle represents the value of another metric. For example, one heat map shows the total memory and percentage of memory use for each virtual machine. Larger rectangles are virtual machines with more total memory, green indicates low memory use, and red indicates high use.

vRealize Operations Manager updates the heat maps in real time as new values are collected for each object and metric. The colored bar below the heat map is the legend. The legend identifies the values that the endpoints represent and the midpoint of the color range.

Heat map objects are grouped by parent. For example, a heat map that shows virtual machine performance, groups the virtual machines by the ESX hosts on which they run.

Create a Custom Heat Map

You can define an unlimited number of custom heat maps to analyze exactly the metrics that you need.

Procedure

1. In the menu, click **Environment**.
2. Select an object to inspect from an inventory tree.
3. Click the **Heat Maps** tab under the **Details** tab.
4. Select the tag to use for first-level grouping of the objects from the **Group By** drop-down menu.

   If a selected object does not have a value for this tag, it appears in a group called Other Groups.
5. Select the tag to use to separate the objects into subgroups from the **Then By** drop-down menu.

   If a selected object does not have a value for this tag, it appears in a subgroup called Other Groups.
6. Select a **Mode** option.

Option	Description
Instance	Track all instances of a metric for an object with a separate rectangle for each metric.
General	Pick an specific instance of a metric for each object and track only that metric.

7 If you selected General mode, select the attribute to use to set the size of the rectangle for each resource in the Size By list and the attribute to use to determine the color of the rectangle for each object in the Color By list.

   Objects that have higher values for the Size By attribute have larger areas in the heat map display. You can also select fixed-size rectangles. The color varies between the colors you set based on the value of the Color By attribute.

   In most cases, the attribute lists include only metrics that vRealize Operations Manager generates. If you select an object type, the list shows all of the attributes that are defined for that object type.

   a To track metrics only for objects of a particular kind, select the object type from the **Object Type** drop-down menu.

8 If you selected Instance mode, select an attribute kind from the **Attribute Kind** list.

   The attribute kind determines the color of the rectangle for each object.

9 Configure colors for the heat map.

   a Click each of the small blocks under the color bar to set the color for low, middle, and high values.

      The bar shows the color range for intermediate values. You can also set the values to match the high and low end of the color range.

   b (Optional) Enter 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.

10 Click **Save** to save the configuration.

   The custom heat map you created appears in the list of heat maps on the **Heat Maps** tab.

Find the Best or Worst Performing Objects for a Metric

You can use heat maps to find the objects with the highest or lowest values for a particular metric.

Prerequisites

If the combination of metrics that you want to compare is not available in the list of defined heat maps, you must define a custom heat map first. See Create a Custom Heat Map.

Procedure

1 In the menu, click **Environment** and select an object from an inventory tree.

2 Click the **Heat Map**s tab under the **Details** tab.

All metric heat maps related to the selected resource appear in the list of predefined heat maps.

3 In the list of heat maps, click the map to view.

The name and metrics values for each object shown on the heat map appear in the list below the heat map.

4 Click the column header for the metric you are interested in to change the sort order, so that the best or worst performing objects appear at the top of the column.

Compare Available Resources to Balance the Load Across the Infrastructure

A heat map can be used to compare the performance of selected metrics across the virtual infrastructure. You can use this information to balance the load across ESX hosts and virtual machines.

Prerequisites

If the combination of metrics to compare is not available in the list of defined heat maps, you must define a custom heat map first. See Create a Custom Heat Map.

Procedure

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

2 Select an object to inspect from an inventory tree.

3 Click the **Heat Maps** tab under the **Details** tab.

4 In the list of heat maps, click the one to view.

The heat map of the selected metrics appears, sized and grouped according to your selection.

5 Use the heat map to compare objects and click resources and metric values for all objects in your virtual environment.

The list of names and metric values for all objects shown on the heat map appear in the list below the heat map. You can click column headers to sort the list by column. If you sort the list by a metric column, you can see the highest or lowest values for that metric on top.

6 (Optional) To see more information about an object in the heat map, click the rectangle that represents this object or click the pop-up window for more details.

What to do next

Based on your findings, you can reorganize the objects in your virtual environment to balance the load between ESX hosts, clusters or datastores.

Heat Maps Tab

With the vRealize Operations Manager heat map feature, you can locate trouble areas based on the metric values for objects in your virtual infrastructure. vRealize Operations Manager uses

analytics algorithms that you can use to compare the performance of objects across the virtual infrastructure in real time using heat maps

How Heat Maps Work

You can use predefined heat maps or create your own custom heat maps to compare the metric values of objects in your virtual environment. vRealize Operations Manager has predefined heat maps on the Details tab that you can use to compare commonly used metrics

Where You Find Heat Maps

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Details** tab, then select the **Heat Maps** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

The **Heat Maps** tab is divided into two panels and the heat map appears between the panels. In the top panel you can create, edit, delete, or clone heat maps. The heap map display depends on the object you select from the environment and the heat map you select.

Table 5-50. Heat Map List Table Columns

Column	Description
Name	Name of the heat map.
Group By	First-level grouping of the objects in the heat map.
Color By	Determines the color of the rectangle for each object.
Size By	An attribute to set the size of the rectangle for each object.
Object Type	Type of object.

The bottom panel updates, depending on what you select on the top panel. In the bottom panel of the **Heat Map** tab you can see the data of the object, calculated by a selected view from the top panel. For example, if the selected object is a host, the result is a list of all the objects on that host.

The Heat Map Display

A heat map displays rectangles of different sizes and colors, and each rectangle represents an object in your virtual environment. The color of the rectangle represents the value of one metric, and the size of the rectangle represents the value of another metric.

vRealize Operations Manager updates the heat maps in real time as new values are collected for each object and metric. The colored bar below the heat map is the legend. The legend identifies the values that the endpoints represent and the midpoint of the color range.

Click a link in the pop-up window for an object to see more details.

Heat Map Configuration Options Workspace

If no predefined heat map shows the information you want to see, you can define a custom heat map. You can select the objects and metrics it tracks, the colors it uses, and the end points for its value range.

Where You Find the Heat Map Configuration Workspace

Select **Environment** in the left pane and select an object from an inventory tree. On the **Details** tab, select **Heat Maps**. On the **Heat Maps** tab, click the plus sign to create a custom heat map.

Table 5-51. Heat Map Configuration Options

Option	Description	
Configurations	■ Add a new configuration. ■ Edit a custom configuration. ■ Delete selected configuration. ■ Clone selected configuration.	
Description	Meaningful description of the heat map.	
Group by	First-level grouping of the objects in the heat map.	
Then by	Subgroups of the first-level object groups in the heat map.	
Mode	General Mode	The heat map shows a colored rectangle for each selected object. 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 an object. 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 object kind.
Size by	Attribute to set the size of the rectangle for each object. Objects that have higher values for the Size by attribute have larger areas of the heat map 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 an object kind, the list shows all of the attributes that are defined for the object type.	
Color by	Determines the color of the rectangle for each object.	
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.	

Using Heat Maps to Analyze Data for Capacity Risk

Planning for capacity risk involves analyzing data to determine how much capacity is available and whether you make efficient use of the infrastructure.

Identify Clusters That Have Enough Space for Virtual Machines

Identify the clusters in a datacenter that have enough space for your next set of virtual machines.

Procedure

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

2 Select **vSphere World**.

3 Click the **Heat Map** tab under the **Details** tab.

4 Select the **Which clusters have the most free capacity and least stress?** heat map.

5 In the heat map, point to each cluster area to view the percentage of remaining capacity.

A color other than green indicates a potential problem.

6 Click **Details** in the pop-up window to examine the resources for the cluster or datacenter.

What to do next

Identify the green clusters with the most capacity to store virtual machines.

Examine Abnormal Host Health

Identifying the source of a performance problem with a host involves examining its workload.

Procedure

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

2 Select **vSphere World**.

3 Click the **Heat Map** tab under the **Details** tab.

4 Select the **Which hosts currently have the most abnormal workload?** heat map.

5 In the heat map, point to the cluster area to view the percentage of remaining capacity.

A color other than green indicates a potential problem.

6 Click **Details** for the ESX host in the pop-up window to examine the resources for the host.

What to do next

Adjust workloads to balance resources as necessary.

Identify Datastores with Enough Space for Virtual Machines

Identify the datastores that have the most space for your next set of virtual machines.

Procedure

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

2 Select **vSphere World**.

3 Click the **Heat Map** tab under the **Details** tab.

4 Select the **Which datastores have the highest disk space overcommitment and the lowest time remaining?** heat map.

5 In the heat map, point to each datacenter area to view the space statistics.

6 If a color other than green indicates a potential problem, click **Details** in the pop-up window to investigate the disk space and disk I/O resources.

What to do next

Identify the datastores with the largest amount of available space for virtual machines.

Identify Datastores with Wasted Space

To improve the efficiency of your virtual infrastructure, identify datastores with the highest amount of wasted space that you can reclaim .

Procedure

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

2 Select **vSphere World**.

3 Click the **Heat Map** tab under the **Details** tab.

4 Select the **Which datastores have the most wasted space and total space storage?** heat map.

5 In the heat map, point to each datacenter area to view the waste statistics.

6 If a color other than green indicates a potential problem, click **Details** in the pop-up window to investigate the disk space and disk I/O resources.

What to do next

Identify the red, orange, or yellow datastores with the highest amount of wasted space.

Identify the Virtual Machines with Resource Waste Across Datastores

Identify the virtual machines that waste resources because of idle, oversized, or powered-off virtual machine states or because of snapshots.

Procedure

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

2 Select **vSphere World**.

3 Click the **Heat Map** tab under the **Details** tab.

4 Select the **For each datastore, which VMs have the most wasted disk space?** heat map.

5 In the heat map, point to each virtual machine to view the waste statistics.

6 If a color other than green indicates a potential problem, click **Details** for the virtual machine in the pop-up window and investigate the disk space and I/O resources.

What to do next

Identify the red, orange, or yellow virtual machines with the highest amount of wasted space.

Examining Relationships in Your Environment

Most objects in an environment are related to other objects in that environment. The **Environment** tab shows how objects in your environment are related. You use this display to troubleshoot problems that might not be about the object that you originally chose to examine. For example, a problem alert on a host might be because a virtual machine related to the host lacks capacity.

Environment Tab Selections

When you select an object from the inventory of your environment, you can display the related objects in an overview, list, or map.

- The Overview shows all the objects in your environment with a status badge for each object. By clicking a badge, you can see which objects are related.
- The List shows only the objects related to your object selection. Depending on the object selected, you can initiate an action or launch an external application.
- The Map shows the objects as icons in a hierarchical display. You select an icon to display the number of related objects.

Use the Overview to identify objects in your environment with health, risk, or efficiency problems. Depending on the object type, you might be able to take action on the object from the List view.

Use the Environment Overview to Find Problems

If you are system administrator who is trying to investigate the reason for slow performance in your environment, you can select key objects such as host systems to see if any related objects such as virtual machines indicate problems.

Procedure

1. In the menu, click **Environment**, then click **vSphere Hosts and Clusters** in the left pane and select the **vSphere World** object.
2. Select the **Environment** tab.

   The system displays health badges for all objects in the vSphere World.
3. Click each of the host system badges.

   The health badge of the virtual machines that belong to the host are highlighted. A host that displays a good health badge, may have virtual machines that display a warning status.

What to do next

Investigate the reason for the problem. For example, once you determine if the problem is chronic or temporary, you can decide how to address it. See Using Troubleshooting Tools to Resolve Problems.

Environment Objects Overview Tab

vRealize Operations Manager collects data for all objects in your environment. You can compare the status of an object with the status of all related objects to determine the possible cause for a problem in your environment.

How the Environment Objects Overview Works

When you select an object in your inventory, vRealize Operations Manager highlights badges for the object and all its related objects. Point to a badge to display current key conditions for an object. See Analyzing the Resources in Your Environment.

Where You Find the Environment Objects Overview

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Environment** tab, then select the **Overview** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 5-52. Environment Objects Overview Options

Option	Description
Badge	Displays the selected badge with the color appropriate to the state of the badge.
Status	All statuses appear by default. Select a status to toggle off the display of badges.
Power State Options	Toggle on to display badges for objects in the on, off, standby, or unknown power states. Selections are additive. For example, you can display objects in both the on and off states. Actions depend on the power state of the object. Use the display to help determine why an action for an object might not be available. See List of vRealize Operations Manager Actions.
Sort	Changes the order in which the objects are listed. Alphabetical sort is by object name.

Environment Objects List Tab

When you select an object in your inventory and choose the list view, vRealize Operations Manager lists all the objects related to your selection. From that list, you can select an object on which to perform an action or to link to another application for information about the object.

Where You Find the Environment Objects List

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Environment** tab, then select the **List** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Each related object is listed with badges that display the status of current key conditions for the object. See Analyzing the Resources in Your Environment.

Point to a badge and click to display a sparkline chart of the condition over time for the object. The chart is qualitative. If a chart shows an inconsistent condition, you might want to investigate events that have occurred for that object. See Events Tab.

Table 5-53. Environment Objects List 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.

Environment Objects Map Tab

When you select an object in your inventory and choose the map view, vRealize Operations Manager displays icons for all the objects related to your selection in a hierarchy. Use the map to see how the objects are related and to get details on any of the objects displayed.

Where You Find the Environment Objects Map

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Environment** tab, then select the **Map** button.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Click any object icon to display the related object types and quantities of each. Depending on your environment, the map display can be very large. Use the map options to control the display of the hierarchy.

Table 5-54. Environment Objects Map Options

Option	Description
Reset Zoom	Resets the map display to fit in the available space.
Pan (no icon)	Click and drag the display to show different parts of the map.
Show values on point (no icon)	Point the mouse to an object icon in the map to display the object name or IP address, and the object type. To display the child object icons for an object in the map, in the pop-up menu that appears, click **Details**. To display the related object types and the number of related objects for an object in the map, click the object icon in the map.
Zoom in	Zooms in on the map. Can also be accomplished using the mouse wheel.
Zoom out	Zooms out on the map. Can also be accomplished using the mouse wheel.
Reset to initial object	If you changed the display to zoom, pan, or select another object, click this option to return to the original display of the initial object.
Object detail	If selected, the display of related object types and quantities for each is limited to the object types that are below the selected object n the hierarchy.

User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options

One of your customers reports poor performance for his virtual machine, including slowness and fails. This scenario provides one way that you can use vRealize Operations Manager to investigate the problem based on information available in the **Troubleshooting** tabs.

As a virtual infrastructure administrator, you respond to a help ticket in which one of your customers reports problems with his virtual machine, sales-10-dk. The reported conditions are poor application performance, including slow load times and slow boot, some of his programs are taking longer and longer to load, and his files are taking longer to save. Today his programs started to fail and an update failed to install.

When you look at the **Alerts** tab for the virtual machine you see an alert for chronic high memory workload leading to memory stress, where the triggered symptoms indicate memory stress and the recommendation is to add memory.

Based on past experience, you are not convinced that this alert indicates the root cause, so you review the **Analysis** tabs. All of the associated badges are green except for Capacity Remaining, which indicates memory and disk space problems, and Time Remaining, which has 0 days remaining for memory and disk space.

From this initial review, you know that problems exist in addition to the memory alert, so you use the **Events** tabs to do a more thorough investigation.

Review the Triggered Symptoms When You Troubleshoot a Virtual Machine Problem

As a virtual infrastructure administrator, you respond to customer complaints and alerts, and identify problems that occur on the objects in your environment. You use the information on the **Symptoms** tab to help determine whether the triggered symptoms indicate conditions that contribute to the reported or identified problem.

You must research a problem of poor performance on one of your virtual machines, as reported by one of your customers. When you view the **Alerts** tab for the virtual machine, the only alert that appears is named `Virtual Machine is Violating Risk Profile 1 in vSphere Hardening Guide`.

When you reviewed the **Analysis** tabs for the virtual machine, you identified that problems were occurring with memory and disk space. Now, you focus your attention to the triggered symptoms on the virtual machine.

The following method of using the **Symptoms** tab to evaluate problems is provided as an example for using vRealize Operations Manager, and is not definitive. Your troubleshooting skills and your knowledge of the particular aspects of your environment determine which methods work for you.

Procedure

1 In the menu, click **Dashboards**, then click **Troubleshoot a VM** in the left pane, .

2 Search for a virtual machine to troubleshoot.

   In this example, the virtual machine name is named **sales-10-dk**.

3 With the virtual machine selected, click the **Alerts** tab, and click the **Symptoms** tab.

4 Review and evaluate the triggered symptoms.

Option	Evaluation Process
Symptom	Are any of the triggered symptoms related to the critical states you see for memory or disk space?
Status	Are the symptoms active or inactive? Even inactive symptoms can provide information about the past state of the object. To add any inactive symptoms, click **Status: Active** on the toolbar to remove the filter.
Created On	When did the symptoms trigger? How does the time of the triggered symptom compare with the other symptoms?
Information	Can you identify a correlation between the triggered symptoms and the state of the Time Remaining and Capacity Remaining badges?

Results

From your review, you determine that some of the triggered symptoms are associated with compliance alerts for the virtual machine as defined in the *vSphere Hardening Guide*. The violated symptoms triggered for the alert named *vSphere Hardening Guide*, which is one of several compliance risk profiles provided with vRealize Operations Manager.

The following symptoms triggered in the compliance alert named `Virtual Machine is Violating Risk Profile 1 in vSphere Hardening Guide`:

- `Independent nonpersistent disks are being used`
- `Autologon feature is enabled`
- `Copy/paste operations are enabled`
- `Users and processes without privileges can remove, connect and modify devices`
- `Guests can receive host information`

Other symptoms also triggered, which are related to memory and time remaining.

- `Guest file system overall disk space usage reaching critical limit`
- `Virtual machine disk space time remaining is low`
- `Virtual machine CPU time remaining is low`
- `Guest partition disk space usage`
- `Virtual machine memory time remaining is low`

What to do next

Review the symptoms for the object on a timeline. See Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem.

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

Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem

Looking at the triggered symptoms for an object over time enables you to compare triggered symptoms, alerts, and events when you are troubleshooting problems with objects in your environment. The **Timeline** tab in vRealize Operations Manager provides a visual chart on which to see triggered symptoms that you can use to investigate problems in your environment.

After you identify the following symptoms as possible indicators of the root cause of the reported performance problems on the sales-10-dk virtual machine, you compare them to each other over time. Look for unusual or common patterns.

- Guest file system overall disk space use reaching critical limit
- Virtual machine disk space time remaining low
- Virtual machine CPU time remaining low
- Guest partition disk space use
- Virtual machine memory time remaining is low

The following method of evaluating problems using the **Timeline** tab is provided as an example for using vRealize Operations Manager and only one method. Your troubleshooting skills and your knowledge of the specifics of your environment determine which methods work for you.

Prerequisites

Review the triggered object symptoms. See Review the Triggered Symptoms When You Troubleshoot a Virtual Machine Problem.

Procedure

1. Enter the name of the virtual machine in the **Search** text box on the main title bar.

   In this example, the virtual machine name is **`sales-10-dk`**.

2. Click the **Events** tab and click the **Timeline** tab.

3. On the Timeline toolbar, click **Date Controls** and select a time that is on or before the reference symptoms were triggered.

   The default time range is the last 6 hours. For a broader view of the virtual machine over time, configure a range that includes triggered symptoms and generated alerts.

4. To view the point at which the symptoms were triggered and to identify which line represents which symptom, drag the timeline week, day, or hour section left and right across the page.

5. Click **Event Filters** and select all the event types.

   Consider whether events correspond to triggered symptoms or generated alerts.

6 In the Related Hierarchies list in the upper left pane, click **vSphere Hosts and Clusters**.

The available ancestors and descendant objects depend on the selected hierarchy.

7 To see if the host is experiencing a contributing problems, click **View From** and select **Host System** under Parent.

Consider whether the host has symptoms, alerts, or events that provide you with more information about memory or disk space problems.

Results

Comparing virtual machine symptoms to host symptoms, and looking at the symptoms over time indicates the following trends:

- The host resource use, host disk use, and host CPU use symptoms are triggered for about 10 minutes approximately every 4 hours.
- The virtual machine guest file system out of space symptom is triggered and canceled over time. Sometimes the symptom is active for an hour and canceled. Sometimes it is active for two hours. But no more than 30 minutes occur between cancellation and the next triggering of the symptom.

What to do next

Look at events in the context of the analysis badges and alerts. See Identify Influential Events When You Troubleshoot a Virtual Machine Problem.

Identify Influential Events When You Troubleshoot a Virtual Machine Problem

Events are changes to objects in your environment that are based on changes to metrics, properties, or information about the object. Examining the events for the problematic virtual machine in the context of the analysis badges and alerts might provide visual clues to the root cause of a problem.

As a virtual infrastructure administrator investigating a reported performance problem with a virtual machine, you compared symptoms on the timeline and identified interesting behavior around the guest file system that you want to examine in the context of other badge metrics to determine if you can find the root cause of the problem.

The following method of evaluating problems using the **Events** tab is provided as an example for using vRealize Operations Manager and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

Prerequisites

Examine triggered symptoms, alerts, and events over time. See Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem

Procedure

1 Enter the name of the virtual machine in the **Search** text box, located on the main title bar.

   In this example, the virtual machine name is sales-10-dk.

2 Click the **Events** tab and select the **Events** button.

3 On the Events toolbar, click **Date Controls** and select a time that is on or before the symptoms were triggered.

4 Click **Event Filters** and select all of the event types.

   Consider whether any changes correspond to other events.

5 Click **View From > Parent > > Select All** and click through the badges in the timeline to review events.

   Consider whether any of the events, which are listed in the data grid below the chart, correspond to problems with the host that might contribute to the reported problem.

6 Click **View From > Child > > Select All** and click through the badges on the toolbar to review the events.

   Consider whether any of the events show problems with the datastore.

Results

Your evaluation shows no particular correlation between the workload or anomalies and the time at which the guest file system out of space symptom was triggered each time.

Running Actions from vRealize Operations Manager

The actions available in vRealize Operations Manager allow you to modify the state or configuration of selected objects in vCenter Server from vRealize Operations Manager. For example, you might need to modify the configuration of an object to address a problematic resource issue or to redistribute resources to optimize your virtual infrastructure.

The most common use of the actions is to solve problems. You can run them as part of your troubleshooting procedures or add them as a resolution recommendation for alerts.

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, and not only on objects that the user can access outside of vRealize Operations Manager.

When you are troubleshooting problems, you can run the actions from the center pane Actions menu or from the toolbar on list views that contain the supported objects.

When an alert is triggered, and you determine that the recommended action is the most likely way to resolve the problem, you can run the action on one or more objects.

Run Actions From Toolbars in vRealize Operations Manager

When you run actions in vRealize Operations Manager, you change the state of vCenter Server objects. You run one or more actions when you encounter objects where the configuration or state of the object is affecting your environment. These actions allow you to reclaim wasted space, adjust memory, or conserve resources.

This procedure for running actions is based on the vRealize Operations Manager **Actions** menus and is commonly used when you are troubleshooting problems. The available actions depend on the type of objects with which you are working. You can also run actions as alert recommendations.

Prerequisites

- Verify that the vCenter Adapter is configured to run actions for each vCenter Server instance. See Configure a vCenter Adapter Instance in vRealize Operations Manager
- Ensure that you understand how to use the Power Off Allowed option if you are running Set CPU Count, Set Memory, and Set CPU Count and Memory actions. See Working With Actions That Use Power Off Allowed.

Procedure

1. Select the object in the Environment page inventory trees or select one or more objects it in a list view.
2. Click **Actions** on the main toolbar or in an embedded view.
3. Select one of the actions.

   If you are working with a virtual machine, only the virtual machine is included in the dialog box. If you are working with clusters, hosts, or datastores, the dialog box that appears includes all objects.
4. Select the check box to run the action on the object, and click **OK**.

   The action runs and a dialog box appears that displays the task ID.
5. To view the status of the job and verify that the job finished, click **Recent Tasks** or click **OK** to close the dialog box.

   The Recent Tasks list appears, which includes the task you just started.

What to do next

To verify that the job completed, click **Environment** in the menu and click **History >Recent Tasks**. Find the task name or task ID in the list and verify that the status is finished. See Monitor Recent Task Status.

Rebalance Container Action

When the workload in your environment becomes imbalanced, you can move the workload across your objects to rebalance the overall workload. The container for the rebalance action can

be a data center or a custom data center, and the objects that are moved are the virtual machines in the recommended list provided by the action.

DRS Must be Enabled on Clusters

Your vCenter Server instance must have a cluster that passes a DRS enabled check for the Rebalance Container action to appear in the Actions drop-down menu.

To get the Rebalance Container action from a custom datacenter or datacenter, and the related alerts, you must have the following:

- A vCenter Adapter configured with the actions enabled for each vCenter Server instance
- A vCenter Server instance with at least one cluster that is DRS enabled

If your cluster does not have DRS fully automated, the Rebalance Container action notifies you that one or more clusters under the selected container do not have DRS set to fully automated.

To ensure that the Rebalance Container action is available in your environment, you must add DRS. Then, wait one collection cycle for the Rebalance Container action to appear.

You Must Have Access to All Objects in the Container

If you have access to all objects in a cluster, data center, or custom data center, you can run the Rebalance Container action to move virtual machines to other clusters. When you do not have access to all of the objects in the container, the Rebalance Container action is not available.

How the Rebalance Container Action Works

If one data center in your environment is experiencing a high workload, while another data center in the same environment is experiencing a low workload, you can use the Rebalance Container action to balance the workload across those objects. For example, if the CPU demand on a host in one data center exceeds its available CPU capacity, critical stress occurs on the host. To identify the cause of stress, monitor the CPU demand. Some virtual machines on each host might be experiencing high CPU demand, whereas others might be experiencing a low demand.

The Rebalance Container action moves all affected objects in the recommended list provided by the action to balance the workload. If you do not want to take action on the entire set of objects to resolve the problem with workload or stress, you can use the Move VM action to move an individual object.

Important Do not attempt to move virtual machines that are members of a vApp, or else the vApp could become nonfunctional. Instead, add affinity rules for these virtual machines to keep them together so that the Move VM and Rebalance Container actions will ignore them.

When workloads become imbalanced, the following alerts can trigger on data centers and custom data centers. These alerts are disabled by default in the policies.

- Custom datacenter has unbalanced workload
- Datacenter has unbalanced workload

When the workloads on hosts in a data center or custom data center differ significantly, click **Home > Alerts** and verify whether the alert triggered. For example, to verify whether the alert triggered on a custom data center, check the alert named `Custom datacenter has unbalanced workload`. You can click the alert to view the causes of the alert and identify the source of the imbalance problem on the **Summary** tab.

To display the recommendations about the objects to move so that you can rebalance the workload, click the **Rebalance Container** action on the **Summary** tab. The recommendations indicate that you move one or more virtual machines to another host. When you click **OK**, a pop-up message provides a link to track the status of the action in **Recent Tasks**.

The action moves the virtual machines identified in the recommendation to the host machine that has a low workload or stress. You can view the status of the action in the list of recent tasks in **Administration > Recent Tasks**. You can also use the vSphere Web Client to view the status of the action and the performance for the host.

After the action runs and vRealize Operations Manager performs several collection cycles, you can view the workload on the data center or custom data center to confirm that the workload was rebalanced and that the alert is no longer triggered.

To see how the workload changed on one or more of your hosts, click a host in the navigation tree. Click **Analysis > Stress** to view the stress score, breakdown, and workload on the host. Then, click **Analysis > Capacity Remaining** to determine how much capacity remains on the host.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab, click **Views**, and select a view of type List.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Recommendations

Review the following information about the hosts and virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Virtual Machine	Name of the virtual machine on the host that is experiencing an excessive workload.

Option	Description
Source Cluster	Name of the cluster on which the virtual machine is running
Datastores	Datastore associated with the virtual machine.
Destination Cluster	Cluster where the virtual machine is to be moved. DRS will select the host automatically.
Reason	Describes the action to be taken and the reason why the move is recommended. For example, the recommendation is to move part of the workload on the cluster to another cluster to reduce the imbalance in CPU demand.
Parent vCenter	Identifies the vCenter vCenter Serveradapter associated with the affected cluster.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-55. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Delete Idle VM Action

The Delete Idle VM action in vRealize Operations Manager removes from your vCenter Server instances selected virtual machines that are in an idle state. Use this action to reclaim redundant resources.

How the Action Works

The Delete Idle VM action removes from your vCenter Server instances virtual machines that are powered on, but that are in idle state.

Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Menu Items

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Menu Items	Description
Name	Name of the virtual machine as it appears in the environment inventory.
Host	Name of the host on which the virtual machine is running.
Parent vCenter	Parent vCenter Server instance where the virtual machine resides.

After you click **Begin Action**, the next dialog box provides the task ID and a link to the task list.

Table 5-56. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set DRS Automation Action

You can monitor and configure the vSphere Distributed Resource Scheduler (DRS) automation rules from vRealize Operations Manager. DRS monitors and allocates the resources in your environment, and balances the computing capacity across your hosts and virtual machines.

How the Action Works

The Set DRS Automation action monitors and configures DRS automation rules. With the Set DRS Automation action, you can enable and disable DRS.

If vRealize Automation manages any of the virtual machines in your environment, the Set DRS Automation action is not available for that object.

Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Menu Items

To ensure that you are submitting the correct action for the correct objects, review the following information about the clusters.

Menu Items	Description
Name	Name of the cluster in the vCenter Server instance.
Automation Level	Level of DRS automation. When DRS is fully automated on the selected cluster, you can run the Set DRS Automation action.
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.
Parent vCenter	Parent vCenter Server instance where the cluster resides.

After you click **Begin Action**, the next dialog box provides the task ID and a link to the task list.

Table 5-57. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Move Virtual Machine Action

You can use the Move VM action to move virtual machines from one host and datastore to another host and datastore to balance the workload in your environment.

How the Action Works

When you initiate this action, the **Move VM** wizard opens and scopes the possible destinations. You select the destination host and datastore from the list of available destinations.

To see all destinations, you must have view access to the following object types:

- Scope object, which includes a vCenter Server, data center, custom data center, or cluster
- Host in the scope object
- Datastore in the host

The destinations include combinations of objects for the move, such as a specific host and datastore, or a different host with the same datastore. You select one of the available combinations. If your environment includes a large number of destination objects, such as many hosts or datastores, enter text in the filter text box to search for specific destination objects.

vRealize Operations Manager uses vSphere DRS rules that you define in vCenter Server to help determine good placement decisions for your virtual machines in the move action. The Affinity Rules column indicates whether those rules will be violated by the Move VM action.

Important Do not attempt to move virtual machines that are members of a vApp, or else the vApp could become nonfunctional. Instead, add affinity rules for these virtual machines to keep them together so that the Move VM and Rebalance Container actions will ignore them.

To initiate the action, you click the **Begin Action** button.

When you finish the wizard, vRealize Operations Manager displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Administration > Recent Tasks**.

Moving Virtual Machines is Not Allowed Across Data Centers

When you attempt to use the **Move VM** action to move a virtual machine across data centers, vRealize Operations Manager must be able to identify the matching network and storage objects for the destination data center. Network objects include VMware virtual switches and distributed virtual switches. Storage objects include datastores and datastore clusters.

Moving a virtual machine across data centers would require vRealize Operations Manager to move the virtual machine files and change the virtual machine network configuration. vRealize Operations Manager does not currently move the virtual machine files across datastores, nor does it change the virtual machine network configuration. Consequently, vRealize Operations Manager does not allow you to move virtual machines across data centers.

When you use the **Move VM** action, be aware of the following behavior:

- If you select a single virtual machine, vRealize Operations Manager displays the data center where the virtual machine resides.
- If you select multiple virtual machines, but those virtual machines do not share a common data center, the **Move VM** action does not display the data centers, and the **Move VM** action does not appear in the actions menu.

Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.

- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Priority	Indicates the priority of the proposed move destination. When the action is automated, the proposed destination with priority of 1 is automatically selected.
Destination Host	Name of the host to which the virtual machine will be moved.
Current CPU Workload	Amount of CPU in GHz available on the host.
Current Memory Workload	Amount of memory in GB available on the host.
Destination Datastore	Datastore to which the virtual machines storage will be moved.
Current Disk Space Workload	Amount of disk space available on the datastore.
Will it fit	Calculated estimation of whether the virtual machine will fit on the selected destination.
VM Power Off Required	When set to `No`, the action does not power off the virtual machine before the move. When set to `Yes`, the action powers off the virtual machine before the move takes place, and powers on the virtual machine after the move is complete. If VMware Tools are installed, a guest OS shutdown is used to power off the virtual machine.
Affinity Rules	Indicates whether vSphere DRS rules exist, as defined in vCenter Server. For example, a rule might exist to keep virtual machines together, and another rule might exist to separate virtual machines. This column indicates the following status. ■ Empty. vSphere DRS rules are not defined. ■ Green check mark. The move of virtual machines will not violate affinity rules. ■ Red circle with bar. The move of virtual machines will break affinity rules. If you choose to break the affinity rules, you must resolve any problems manually.
Affinity Rule Details	Identifies the virtual machine and the vSphere DRS rule name as defined in vCenter Server, which will be broken if you move the virtual machine.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-58. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Power Off Virtual Machine Action

The Power Off VM action in vRealize Operations Manager stops one or more selected virtual machines that are in a powered on state. You power off a virtual machine when you are managing resources and reclaiming wasted space.

How the Action Works

The Power Off VM action turns off the virtual machine. If VMware Tools is installed and running, the guest operating system is shutdown before the machine is powered off. If VMware Tools is not installed and running, the virtual machine is powered off regardless of the state of the guest operating system. In this case, use this action only when you are powering off virtual machines where stopping the guest operating system will not adversely affect the installed applications.

If the target virtual machine is already powered off, the recent task status reports success on the machine, even though the state of the virtual machine did not change.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
Name	Name of the virtual machine as it appears in the environment inventory.
Power State	Indicates whether the virtual machine is powered on or powered off.

Option	Description
Idle VM	Indicates whether the virtual machine is considered to be in the idle state based on the configured idle virtual machine metric. Possible values include: ■ false. The virtual machine is active. ■ true. The virtual machine is idle. ■ unknown. vRealize Operations Manager does not have the data required to calculate the idle metric.
Idle VM Percentage	Calculated threshold of the idle virtual machine percentage based on the configured reclaimable wasted space policy.
CPU Usage Percentage	Calculated threshold of the virtual machine CPU percentage based on the metric named `cpu | usage_average`.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-59. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Shut Down Guest Operating System for Virtual Machine Action

The Shut Down Guest OS for VM action shuts down the guest operating system and powers off the virtual machine. You shut down a virtual machine when you are managing resources and reclaiming wasted space.

How the Action Works

The Shut Down Guest OS for VM action checks that VMware Tools, which is required, is installed on the target virtual machines, then shuts down the guest operating system and powers off the virtual machine. If VMware Tools is not installed or installed but not running, the action does not run and the job is reported as failed in **Recent Tasks**.

If the target virtual machine is already powered off, the recent task status reports success on the machine, even though the state of the virtual machine did not change.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.

- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
Name	Name of the virtual machine as it appears in the environment inventory.
Power State	Indicates whether the virtual machine is powered on or powered off.
Idle VM	Indicates whether the virtual machine is considered to be in the idle state based on the configured idle virtual machine metric. Possible values include: ■ false. The virtual machine is active. ■ true. The virtual machine is idle. ■ unknown. vRealize Operations Manager does not have the data required to calculate the idle metric.
Idle VM Percentage	Calculated threshold of the idle virtual machine percentage based on the configured reclaimable wasted space policy.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-60. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Power On Virtual Machine Action

Use the Power On VM action in vRealize Operations Manager to start one or more virtual machines that are in a powered off state. You power on a virtual machine so that you can shift

resources. For example, power on a machine so that you can use it, run applications, or verify that actions that were run on already powered down machines contribute to improved performance.

How the Action Works

The Power On VM action powers on virtual machines that are powered off. Virtual machines that are currently powered on are not affected by the action.

If the target virtual machine is already powered on, the task status reports success for the machine even though the state of the virtual machine did not change.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
Name	Name of the virtual machine as it appears in the environment inventory.
Power State	Indicates whether the virtual machine is powered on or powered off.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-61. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Delete Powered Off Virtual Machine Action

The Delete Powered Off VM action in vRealize Operations Manager removes selected virtual machines that are in a powered off state from your vCenter Server instances. Use this action to reclaim redundant resources.

How the Action Works

The Delete Powered Off VM action removes virtual machines from the vCenter Server instances. If the virtual machine is powered on, the action does not delete the virtual machine.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
Name	Name of the virtual machine as it appears in the environment inventory.
Power State	Indicates whether the virtual machine is powered on or powered off.

Option	Description
Disk Space	Amount of disk space currently consumed by the virtual machine.
Snapshot Space	Amount of disk space currently consumed by the virtual machine snapshots.
Memory (MB)	Amount of memory allocated to the virtual machine.
CPU Count	Number of CPUs currently configured for the virtual machine.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-62. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set Memory for Virtual Machine Action

The Set Memory for VM action in vRealize Operations Manager is used to add or remove memory on virtual machines. You increase the memory to address performance problems or decrease the memory to reclaim resources.

How the Action Works

The Set Memory for VM action determines the power state of the target virtual machines, takes a snapshot if you request it, powers off the machine if required and requested, changes the memory to the new value, and returns the virtual machines their original power states.

An alternative form of the Set Memory for Virtual Machine action is available for automation. This action can run when the virtual machine is powered on or off.

Use this version of the action if the automated action has permission to power off the virtual machine, and hot add of memory is not enabled on the virtual machine. With hot add enabled, you can add memory, but you cannot remove it.

This version of the action would be required if a virtual machine is powered on and the amount of memory must be reduced.

This version of the action has the Power Off Allowed flag set to true. You can select this Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is enabled on the virtual machines, then power off is not required. If power off is required and VMware Tools are installed, then the virtual machines are shut down before they are powered off.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. If you modify a value, the check box is selected. The check box must be selected to enable the OK button.
Name	Name of the virtual machine as it appears in the environment inventory.
New (MB)	Requested amount of memory in megabytes. The value must be a multiple of 4, and must not be less than 4. If the value is less than 4 or is not a multiple of 4, the amount of memory does not change, and Recent Tasks displays the action as failed. ■ When the virtual machine power state is `PoweredOn`, the memory hot plug configuration limits of the virtual machine are factored into the requested amount and might result in a different configured memory than requested. ■ If memory hot plug is not enabled, the request fails unless you also select **Power Off Allowed**. ■ If memory hot plug is enabled, the configured memory is adjusted to be a multiple of the virtual machine hot plug memory increment, and at least that increment more than the current virtual machine memory configuration. The adjusted memory configuration must also be no more than the hot plug memory limit. If the memory hot plug constraints of the virtual machine cannot be satisfied, the amount of memory does not change, and Recent Tasks displays the action as failed unless you also you also select **Power Off Allowed**. If Power Off Allowed is selected, the action first attempts to satisfy the memory reconfiguration request without powering off the virtual machine, and only powers off the virtual machine if it is necessary to reconfigure the memory.
Current (MB)	Amount of memory in megabytes that is currently configured on the virtual machine.

Option	Description
Power State	Indicates whether the virtual machine is powered on or powered off.
Power Off Allowed	If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system. In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied. See Working With Actions That Use Power Off Allowed.
Snapshot	Creates a snapshot of the virtual machine before modifying the memory. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results. The name of the snapshot is supplied in the Recent Tasks messages for the action. If the memory is changed with Memory Hot Plug enabled, then the snapshot is taken with the virtual machine is running, which consumes more disk space.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-63. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set Memory Resources for Virtual Machine Action

The Set Memory Resources for VM action is used to modify the memory reservation and memory limit on virtual machines. You modify the memory reservation and limit to manage resources in your environment, either to reclaim unused resources or to ensure that your virtual machines have the resources they need to run efficiently.

How the Action Works

The Set Memory Resources for VM action determines how memory resources are allocated to the virtual machine. The reservation value is the minimum amount of guaranteed memory allocated for the virtual machine. The limit is the maximum amount of memory that the virtual machine can consume.

The reservation and limit values in vCenter Server are set in megabytes. vRealize Operations Manager calculates and reports on memory in kilobytes. When you run this action, the values are presented in kilobytes so that you can implement recommendations from vRealize Operations Manager.

To run the action, all options must be configured in the dialog box for the objects on which your are running the action. If you are changing one option to a new value, but not another option, ensure that the option that you do not want to change is configured with the current value.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. If you modify a value, the check box is selected. The check box must be selected to enable the OK button.
Name	Name of the virtual machine as it appears in the environment inventory.
New Resv (KB)	Amount of memory in kilobytes reserved for the virtual machine when the action is finished. The new reservation value must be less than or equal to the new limit value unless your new limit is unlimited (-1). The reservation supports the following possible values: ■ If you set the value to 0, the virtual machine is allocated only the currently configured amount of RAM. ■ If you add or remove reserved memory, the value must be evenly divisible by 1024.
Current Resv (KB)	Amount of memory in kilobytes that is currently configured as the guaranteed memory for the virtual machine.
New Limit (KB)	Maximum amount of memory in kilobytes that the virtual machine can consume when the action is completed. The limit supports the following possible values: ■ If you set the value to 0, then the maximum memory is no greater than the allocated reservation amount. ■ If you set the value to -1, then the virtual machine memory is unlimited. ■ It you increase or decrease the limit, the value must be evenly divisible by 1024.

Option	Description
Current Limit (KB)	Maximum amount of memory that the virtual machine is currently allowed to consume.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-64. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set CPU Count for Virtual Machine Action

The Set CPU action in vRealize Operations Manager modifies the number of vCPUs on a virtual machine. You increase the number of CPUs to address performance problems or decrease the number of CPU to reclaim resources.

How the Action Works

The Set CPU Count action shuts down or powers off the target virtual machines, which is required if you are decreasing the CPU count. This action creates a snapshot if you request it, changes the number of vCPUs based on the new CPU count you provided, and returns the virtual machines to their original power states.

An alternative form of the Set CPU Count for Virtual Machine action is available for automation. This action can run when the virtual machine is powered on or off.

Use this version of the action if the automated action has permission to power off the virtual machine, and hot add of memory is not enabled on the virtual machine. With hot add enabled, you can add CPUs, but you cannot remove them.

This version of the action would be required if a virtual machine is powered on and the number of CPUs must be reduced.

This version of the action has the Power Off Allowed flag set to true. You can select this Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is enabled on the virtual machines, then power off is not required. If power off is required and VMware Tools are installed, then the virtual machines are shut down before they are powered off.

Set CPU Count for a Virtual Machine

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_set_cpu_count_for_vm)

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. If you modify a value, the check box is selected. The check box must be selected to enable the OK button.
Name	Name of the virtual machine as it appears in the environment inventory.
New CPU	Number of CPUs when the action is completed. If the value is less than 1 or a value not supported for the virtual machine in vCenter Server, and the virtual machine is powered on and Hot Add is not enabled, the number of CPUs does not change and Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action. The value that appears is the calculated recommended size. If the target virtual machine is new or offline, this value is the current number of CPUs. If vRealize Operations Manager has been monitoring the virtual machine for 6 or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.
Current CPU	Number of configured CPUs.
Power State	Indicates whether the virtual machine is powered on or powered off.

Option	Description
Power Off Allowed	If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system. In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied. See Working With Actions That Use Power Off Allowed.
Snapshot	Creates a snapshot before changing the number of CPUs. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results. The name of the snapshot is supplied in the Recent Tasks messages for the action. If the CPU is changed with CPU Hot Plug enabled, then the snapshot is taken with the virtual machine is running, which consumes more disk space.
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-65. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set CPU Resources for Virtual Machine Action

The Set CPU Resources for VM action is used to modify the CPU reservation and CPU limit on virtual machines. You modify the CPU reservation and limit to manage workload demands in your environment.

How the Action Works

The Set CPU Resources for VM action determines how CPU resources can be allocated to the virtual machines. The reservation limit is the minimum amount of guaranteed CPU resources allocated to the virtual machine. The limit is the maximum amount of CPU resources that the virtual machine can consume.

To run the action, all options where you configure a value must contain a value for the objects that you want to change. If you are changing one option to a new value, but not another option, ensure that the option that you are not changing is configured with the current value.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. If you modify a value, the check box is selected. The check box must be selected to enable the OK button.
Name	Name of the virtual machine as it appears in the environment inventory.
New Resv (MHz)	Amount of CPU resources in megahertz reserved for the virtual machine when the action is finished. The new reservation value must be less than or equal to the new limit value unless your new limit is unlimited (-1). The reservation supports the following possible values: ■ If you set the value to 0, the virtual machine is allocated only the currently configured CPU consumption level. ■ If you add or removed reserved CPU consumption, supply a positive integer unless you set the value to 0.
Current Resv (MHz)	Amount of CPU resources that is currently configured as the guaranteed CPU resources for the virtual machine.
New Limit (MHz)	Maximum amount of CPU consumption in megahertz that the virtual machine can consume when the action is completed. The limit supports the following possible values: ■ If you set the value to 0, the maximum CPU consumption is not greater than the allocated reservation amount. ■ If you set the value to -1, then the virtual machine CPU consumption is unlimited. ■ If you add or remove CPU consumption limits, supply a positive integer, unless you set the value to 0 or -1.
Current Limit (MHz)	Maximum amount of CPU that the virtual machine can consume.

Option	Description
Host	Name of the host on which the virtual machine is running.
Adapter Instance	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-66. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Set CPU Count and Memory for Virtual Machine Action

The Set CPU Count and Memory for VM action is used to add or remove CPUs and memory on virtual machines with only one power down of the virtual machines to perform the combined actions. You modify the CPU and memory to address performance problems or to reclaim resources.

How the Action Works

The Set CPU Count and Memory action powers off the target virtual machines, creates a snapshot if requested, changes the number of vCPUs and memory based on the new CPU count and memory values you provided, and returns the virtual machines their original power states.

An alternative form of the Set CPU Count and Memory for Virtual Machine action is available for automation. This version of the action has the Power Off Allowed flag set to true so that the action is available for automation and can run when the virtual machine is in the powered on state. You can select the Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is enabled on the virtual machines, then power off is not required. If power off is required and VMware Tools are installed, then the virtual machines are shut down before they are powered off.

To run the action, all options where you configure a value must contain a value for the objects that you want to change. If you are changing one option to a new value, but not another option, ensure that the option that you are not changing is configured with the current value.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.

- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. If you modify a value, the check box is selected. The check box must be selected to enable the OK button.
Name	Name of the virtual machine as it appears in the environment inventory.
New CPU	Number of CPUs when the action is completed. If the value is less than 1 or a value not supported for the virtual machine in vCenter Server, and the virtual machine is powered on and Hot Add is not enabled, the number of CPUs does not change and Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action. The value that appears is the calculated recommended size. If the target virtual machine is new or offline, this value is the current number of CPUs. If vRealize Operations Manager has been monitoring the virtual machine for 6 or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.
Current CPU	Number of currently configured CPUs.
New (MB)	Amount of memory in megabytes when the action is completed. The value must be a multiple of 4, and not less than 4. If the value is less that 4 or is not a multiple of 4, and the virtual machine is powered on and Hot Add is not enabled, the amount of memory does not change and the Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action. The value that appears is the calculated recommended size. If the target virtual machine is new or offline, this value is the currently configured memory. If vRealize Operations Manager has been monitoring the virtual machine for 6 or more hours, depending on your environment, the value that appears is the Memory Recommended Size metric.
Current (MB)	Amount of memory in megabytes that is currently configured on the virtual machine.
Power State	Indicates whether the virtual machine is powered on or powered off.

Option	Description
Power Off Allowed	If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system. In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied. See Working With Actions That Use Power Off Allowed.
Snapshot	If selected, the action creates a snapshot of the virtual machine before modifying the CPU count and the memory. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results.
Host	Name of the host on which the virtual machine is running.
Adapter	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-67. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

Delete Unused Snapshots for Virtual Machine Action

The Delete Unused Snapshots for Virtual Machines action in vRealize Operations Manager deletes snapshots that are older than the specified age from your datastores. Deleting unused snapshots reclaims wasted space in your environment.

How the Action Works

The Delete Unused Snapshots for Virtual Machine action comprises two dialog boxes. The first dialog box allows you to select the snapshot age criteria, which must be greater than one day. The second step allows you to select the snapshots to delete, and runs the Delete Unused Snapshots for Virtual Machine action.

The number of days that you specify for each virtual machine is the age of the snapshots based on the creation date. The Delete Unused Snapshots for Virtual Machine action retrieves the snapshot and displays the snapshot name, space consumed, and location so that you can evaluate the snapshots before you delete them.

When you click **Begin Action**, vRealize Operations Manager displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Administration > Recent Tasks**.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

You first retrieve snapshots based on age, then select the snapshots to delete.

Table 5-68. Retrieve Snapshots

Option	Description
Name	Name of the virtual machine on which you are running the Delete Unused Snapshots for VM action.
Days Old	Age of the snapshots to be deleted. This action retrieves snapshots for the virtual machine that are older than one day.
Host	Name of the host with which the virtual machine is associated.
Parent vCenter	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

Select the snapshots to delete.

Table 5-69. Delete Snapshots

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
VM Name	Name of the virtual machine from which the snapshot was created.
Snapshot Name	Name of the snapshot in the datastore.
Snapshot Space (MB)	Number of megabytes consumed by the snapshot.

Table 5-69. Delete Snapshots (continued)

Option	Description
Snapshot Create Time	Date and time when the snapshot was created.
Snapshot Age	Age of the snapshot in days.
Datacenter Name	Name of the datacenter with which the datastore is associated.
Datastore Name	Name of the datastore where the snapshot is managed.
Host Name	Name of the host with which the datastore is associated.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-70. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

The Delete Unused Snapshots action creates a job for the retrieve snapshots action, and a job for the delete snapshots action.

Delete Unused Snapshots for Datastore Action

The Delete Unused Snapshots for Datastore action in vRealize Operations Manager deletes snapshots that are older than the specified age from your datastores. Deleting unused snapshots reclaims wasted space in your environment.

How the Action Works

The Delete Unused Snapshots for Datastore action comprises two dialog boxes. The first dialog box allows you to select the snapshot age criteria, which must be greater than one day. The second step allows you to select the snapshots to delete, and runs the Delete Unused Snapshots for Datastore action.

The number of days that you specify for each datastore is the age of the snapshots based on the creation date. The Delete Unused Snapshots dialog box provides details regarding snapshot name, space consumed, and location so that you can evaluate the snapshots before you delete them.

When you click **Begin Action**, vRealize Operations Manager displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Administration > Recent Tasks**.

Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a recommended action on an alert.

For the supported objects and object levels, this action is available in the following locations in vRealize Operations Manager:

- Embedded just below the top menu.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Details** tab and click **Views**.
- On the toolbar when you click **Environment** in the menu, select an object, click the **Environment** tab, and select an object in the list view.
- In the Inventory Explorer list when you click **Administration** in the menu, click **Configuration > Inventory Explorer**, click the **List** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

Action Options

Review the following information about datastores to ensure that you are submitting the action for the correct objects.

You first retrieve snapshots based on age, then select the snapshots to delete.

Table 5-71. Retrieve Snapshots

Option	Description
Name	Name of the datastore on which you are running the delete snapshot action.
Days Old	Age of the snapshots to be deleted. This action retrieves snapshots for the datastore that are older than one day.
Host	Name of the host with which the datastore is associated.
Parent vCenter	Name of the VMware Adapter as it is configured in vRealize Operations Manager. The adapter manages the communication with the vCenter Server instance.

Select the snapshots to delete.

Table 5-72. Delete Snapshots

Option	Description
Selected objects	Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.
Datastore Name	Name of the datastore where the snapshot is managed.
Snapshot Name	Name of the snapshot in the datastore.
Snapshot Space (MB)	Number of megabytes consumed by the snapshot.
Snapshot Create Time	Date and time when the snapshot was created.
Snapshot Age	Age of the snapshot in days.
Datacenter Name	Name of the datacenter with which the datastore is associated.

Table 5-72. Delete Snapshots (continued)

Option	Description
Host Name	Name of the host with which the datastore is associated.
VM Name	Name of the virtual machine from which the snapshot was created.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 5-73. Task ID Dialog Box

Option	Description
Recent Tasks	To view the status of the job and verify that the job finished, click **Recent Tasks**.
OK	To close the dialog box without further action, click **OK**.

The Delete Unused Snapshots action creates a job for the retrieve snapshots action, and a job for the delete snapshots action.

Troubleshoot Actions in vRealize Operations Manager

If you are missing data or cannot run actions from vRealize Operations Manager, review the troubleshooting options.

Verify that your vCenter Adapter is configured to connect to the correct vCenter Server instances, and configured to run actions. See Configure a vCenter Adapter Instance in vRealize Operations Manager.

- Actions Do Not Appear on Object

  An action might not appear on an object, such as a host or virtual machine, because that object is being managed by vRealize Automation.

- Missing Column Data in Actions Dialog Boxes

  Data is missing for one or more objects in an Actions dialog box, making it difficult to determine if you want to run the action.

- Missing Column Data in the Set Memory for VM Dialog Box

  The read-only data columns do not display the current values, which makes it difficult to properly specify a new memory value.

- Host Name Does Not Appear in Action Dialog Box

  When you run an action on a virtual machine, the host name is blank in the action dialog box.

Actions Do Not Appear on Object

An action might not appear on an object, such as a host or virtual machine, because that object is being managed by vRealize Automation.

Problem

Actions such as Rebalance Container might not appear in the drop-down menu when you view the actions for your data center.

- If a data center is managed by vRealize Automation, actions do not appear.
- If a data center is not managed by vRealize Automation, you can take action on the virtual machines that are not being managed by vRealize Automation.

Cause

When vRealize Automation manages the child objects of a data center or custom data center container, the actions that are normally available on those objects do not appear, because the action framework excludes actions on objects that vRealize Automation manages. You cannot turn on or turn off the exclusion of actions on vRealize Automation managed objects. This behavior is normal.

If you removed the vRealize Automation adapter instance, but did not select the **Remove related objects** check box, the actions are still disabled.

To make actions available on the objects in your data center or custom data center, either confirm that vRealize Automation is not managing the objects, or perform the steps in this procedure to remove the vRealize Automation adapter instance.

Solution

1. To allow actions on an object, go to your vRealize Automation instance.
2. Make the change in vRealize Automation, such as to move a virtual machine.

Missing Column Data in Actions Dialog Boxes

Data is missing for one or more objects in an Actions dialog box, making it difficult to determine if you want to run the action.

Problem

When you run an action one or more objects, some of the fields are empty.

Cause

The VMware vSphere adapter has not collected the data from the vCenter Server instance that manages the object or the current vRealize Operations Manager user does not have privileges to view the collected data for the object.

Solution

1. Verify that vRealize Operations Manager is configured to collect the data.
2. Verify that you have the privileges necessary to view the data.

Missing Column Data in the Set Memory for VM Dialog Box

The read-only data columns do not display the current values, which makes it difficult to properly specify a new memory value.

Problem

Current (MB) and Power State columns do not display the current values, which are collected for the managed object.

Cause

The adapter responsible for collecting data from the vCenter Server on which the target virtual machine is running has not run a collection cycle and collected the data. This can occur when you recently created an VMware adapter instance for the target vCenter Server and initiated an action. The VMware vSphere adapter has a 5-minute collection cycle.

Solution

1. After you create a VMware adapter instance, wait an additional 5 minutes.
2. Rerun the **Set Memory for VM** action.

   The current memory value and the current power state appear in the dialog box.

Host Name Does Not Appear in Action Dialog Box

When you run an action on a virtual machine, the host name is blank in the action dialog box.

Problem

When you select virtual machine on which to run an action, and click the **Action** button, the dialog box appears, but the Host column is empty.

Cause

Although your user role is configured to run action on the virtual machines, you do not have a user roll that provides you with access to the host. You can see the virtual machines and run actions on them, but you cannot see the host data for the virtual machines. vRealize Operations Manager cannot retrieve data that you do not have permission to access.

Solution

You can run the action, but you cannot see the host name in the action dialog boxes.

Monitor Recent Task Status

The Recent Task status includes all the tasks initiated from vRealize Operations Manager. You use the task status information to verify that your tasks finished successfully or to determine the current state of tasks.

You can monitor the status of tasks that are started when you run actions, and investigate whether a task finished successfully.

Prerequisites

You ran at least one action as part of an alert recommendation or from one of the toolbars. See Run Actions From Toolbars in vRealize Operations Manager.

Procedure

1. In the menu, click **Administration**, then select **History** from the left pane.
2. Click **Recent Tasks**.
3. To determine if you have tasks that are not finished, click the **Status** column and sort the results.

Option	Description
`In Progress`	Indicates running tasks.
`Completed`	Indicates finished tasks.
`Failed`	Indicates incomplete tasks on at least one object when started on multiple objects.
`Maximum Time Reached`	Indicates timed out tasks.

4. To evaluate a task process, select the task in the list and review the information in the **Details of Task Selected** pane.

   The details appear in the Messages pane. If the information message includes `No action taken`, the task finished because the object was already in the requested state.
5. To view the messages for an object when the task included several objects, select the object in the Associated Objects list.

   To clear the object selection so that you can view all the messages, press the space bar.

What to do next

Troubleshoot tasks with a status of `Maximum Time Reached` or `Failed` to determine why a task did not run successfully. See Troubleshoot Failed Tasks.

Recent Tasks in vRealize Operations Manager

The status of the tasks that were recently initiated from vRealize Operations Manager appears in the Recent Task list. You can determine whether a task is finished, still in process, or failed.

How Recent Tasks Work

The Recent Tasks page reports on logged task events, and the log entries appear in the messages area so that you can troubleshoot failed tasks.

Where You View Recent Tasks

In the menu, select **Administration**, then select **History** from the left pane and click **Recent Tasks**.

Recent Task Options

Review the information in the task list to determine if a task is completed or if you need to troubleshoot a failed task. To see the details about a task, select the task in the list and review the associated objects and task messages.

Table 5-74. Task List

Option	Description
Export	Exports the selected task to an XML file. The exported information, which includes the messages, is useful when you are troubleshooting a problem.
Edit Properties	Determines how long the recent task data is retained in your system. Set the number of days that vRealize Operations Manager keeps the data, after which it is purged from the system. The default value is 90 days.
Status drop-down menu	Filters the list based on the status value.
All Filters	Filters the list based the selected column and the provided values.
Filter (Object Name)	Limits the tasks in the list to those that match the entered string. The search is based on a partial entry. For example, if you enter **vm**, objects such as **vm001** and **acctvm_east** are included.
Task	Name of the task. For example, Set CPU Count for VM.

Table 5-74. Task List (continued)

Option	Description
Status	Current state of the task. Possible states include the following values: ■ Completed. Task completed successfully on the target objects. ■ In Progress. Task is running on the target objects. ■ Failed. Task failed to run on the target objects. If the task started, the reasons for failure might include a faulty script, a script timed out, or actions are not taken. If the task did not start and immediately reports as failed, the reasons might include that the task could not start or the script could was not found. If the task was not initiated on the target object, it might have failed because of communication or authentication errors. ■ Maximum Time Reached. Task is running past the amount of time that is the default or configured value. To determine the current status, you must troubleshoot the initiated action. ■ Not Dispatched. The action adapter was not found. ■ Started. Task is initiated on the object. ■ Unknown. An error occurred while running the action, but the error was not captured in the task logs. To further investigate this status, check the vRealize Operations Manager support logs for the vCenter Adapter, available in the **Administration** area, and check the target system.
Started Time	Date and time when the task started.
Completed Time	Date and time when the task finished. A completed date does not appear if the task failed or if the maximum time out is reached.
Automated	Indicates whether the action in the task list was automated, indicated by `Yes` or `No`.
Object Name	Object on which the task was started.
Object Type	Type of object on which the task was started.
Alert	Alert that triggered the action automatically. When an alert is triggered that is associated to the recommendation, it triggers the action without user intervention. You can automate Alert recommendations that have an associated action. Automation is disabled by default. You configure automation in the **Override Alert / Symptom Definitions** area of a policy when you create or edit the policy in **Administration > Policies**. An administrator who has the **Automation** role has permission to automate actions in the **Override Alert / Symptom Definitions** area of the policy workspace.
Source Type	Authentication source that the user who started the task used when accessing vRealize Operations Manager.

Table 5-74. Task List (continued)

Option	Description
Submitted By	Name of the user who initiated the task. This column displays the automationAdmin user account for automated actions that are triggered by alerts.
Task ID	ID generated when the task, which included one or more actions, was started. The task ID is unique for the task for each adapter. If a task includes tasks that ran using two adapters, you see two task IDs. If the task is a delete snapshot action, two task IDs are generated. One ID is for the retrieve snapshots based on date task, and the other ID is for the delete selected snapshots task.

The Associated Objects are the objects on which the selected task ran.

Table 5-75. Associated Objects for Selected Task Details

Option	Description
Object Name	Detailed list of objects that are included in the task selected in the task list. If the task ran on only one object, the list includes one object. If the task ran on multiple objects, each object is listed on a separate row.
Object Type	Type of object for each object name.
Status	Current state of the task.

The Messages are the log of the task as it ran. You use the logs to identify problems if the task does not finish successfully.

Table 5-76. Messages for Selected Task Details

Severity drop-down menu	Limits the messages based on the Severity value.
Filter (Message)	Limits the message in the list to those that match the entered string. The search is based on a partial entry. For example, if you enter **id**, then messages that contain **Task ID** and the phrase **did not complete** are included.
Severity	Message level in the logs. The severity includes the following values: ■ Information. Messages added to logs as the task is processed. ■ Error. Messages generated during a task failure.

Table 5-76. Messages for Selected Task Details (continued)

Time	Date and time the entry was added to the log.
Message	Text of the log entry. Use the information in the message to determine why a task failed, and to begin to troubleshoot and resolve the failure. The messages appear with the most recent entry at the top of the list if you do not sort the columns.

Troubleshoot Failed Tasks

If tasks fail to run in vRealize Operations Manager, review the Recent Tasks page and troubleshoot the task to determine why it failed.

This information is a general procedure for using the information in Recent Tasks to troubleshoot problems identified in the tasks.

- Determine If a Recent Task Failed

  The Recent Tasks provide the status of action tasks initiated from vRealize Operations Manager. If you do not see the expected results, review the tasks to determine if your task failed.

- Troubleshooting Maximum Time Reached Task Status

  An action task has a `Maximum Time Reached` status and you do not know the current status to the task.

- Troubleshooting Set CPU or Set Memory Failed Tasks

  An action task for Set CPU Count or Set Memory for VM has a `Failed` status in the recent task list because power off is not allowed.

- Troubleshooting Set CPU Count or Set Memory with Powered Off Allowed

  A Set CPU Count, Set Memory, or a Set CPU Count and Set Memory action indicates that the action failed in Recent Tasks.

- Troubleshooting Set CPU Count and Memory When Values Not Supported

  If you run the Set CPU Count or Set Memory actions with an unsupported value on a virtual machine, the virtual machine might be left in an unusable state and require you to resolve the problem in vCenter Server.

- Troubleshooting Set CPU Resources or Set Memory Resources When the Value is Not Supported

  If you run the Set CPU Resources action with an unsupported value on a virtual machine, the task fails and an error appears in the Recent Task messages.

- Troubleshooting Set CPU Resources or Set Memory Resources When the Value is Too High

  If you run the Set CPU Resources or Set Memory Resources action with a value that is greater than the value that your vCenter Server instance supports, the task fails and an error appears in the Recent Tasks messages.

- Troubleshooting Set Memory Resources When the Value is Not Evenly Divisible by 1024

  If you run the Set Memory Resources action with a value that cannot convert from kilobytes to megabytes, the task fails and an error appears in the Recent Task messages.

- Troubleshooting Failed Shut Down VM Action Status

  A shut down VM action task has a `Failed` status in the Recent Task list.

- Troubleshooting VMware Tools Not Running for a Shut Down VM Action Status

  A Shut down VM action task has a `Failed` status in the Recent Task list and the Message indicates that VMware Tools were required.

- Troubleshooting Failed Delete Unused Snapshots Action Status

  A Delete Unused Snapshots action task has a `Failed` status in the Recent Task list.

Determine If a Recent Task Failed

The Recent Tasks provide the status of action tasks initiated from vRealize Operations Manager. If you do not see the expected results, review the tasks to determine if your task failed.

Procedure

1. In the menu, click **Administration**, then click **History** in the left pane
2. Click **Recent Tasks**.
3. Select the failed task in the task list.
4. In the Messages list, locate the occurrences of `Script Return Result: Failure` and review the information between this value and `<-- Executing:[script name] on {object type}`.

   `Script Return Result` is the end of action run and `<-- Executing` indicates the beginning. The information provided includes the parameters that are passed, the target object, and unexpected exceptions that you can use to identify the problem.

Troubleshooting Maximum Time Reached Task Status

An action task has a `Maximum Time Reached` status and you do not know the current status to the task.

Problem

The Recent Tasks list indicates that a task had a status of `Maximum Time Reached`.

The task is running past the amount of time that is the default or configured value. To determine the current status, you must troubleshoot the initiated action.

Cause

The task is running past the amount of time that is the default or configured value for one of the following reasons:

- The action is exceptionally long running and did not finish before the threshold timeout was reached.
- The action adapter did not receive a response from the target system before reaching the timeout. The action might have completed successfully, but the completion status was not returned to vRealize Operations Manager.
- The action did not start correctly.
- The action adapter might have an error and be unable to report the status.

Solution

Check the state of the target object to determine whether the action completed successfully. If it did not, continue investigating to find the root cause.

Troubleshooting Set CPU or Set Memory Failed Tasks

An action task for Set CPU Count or Set Memory for VM has a `Failed` status in the recent task list because power off is not allowed.

Problem

The Recent Tasks list indicates that a Set CPU Count, Set Memory, or Set CPU and Memory task has a status of `Failed`. When you evaluate the Messages list for the selected task, you see this message.

```
Unable to perform action. Virtual Machine found
     powered on, power off not allowed
```

When you increase memory or CPU count, you see this message.

```
Virtual Machine found powered on, power off not allowed, if hot add is
     enabled the hotPlugLimit is exceeded
```

Cause

You submitted the action to increase or decrease the CPU or memory value without selecting the **Allow Power Off** option. When you ran the action where a target object is currently powered on and where **Memory Hot Plug** is not enabled for the target object in vCenter Server, the action fails.

Solution

1. Either enable **Memory Hot Plug** on your target virtual machines in vCenter Server or select **Allow Power Off** when you run the Set CPU Count, Set Memory, or Set CPU and Memory actions.

2 Check your hot plug limit in vCenter Server.

Troubleshooting Set CPU Count or Set Memory with Powered Off Allowed

A Set CPU Count, Set Memory, or a Set CPU Count and Set Memory action indicates that the action failed in Recent Tasks.

Problem

When you run an action that changes the CPU count, the memory, or both, the action fails even though you know that the Power Off Allowed was selected, the virtual machine is running, and the VMware Tools are installed and running.

Cause

The virtual machine should shut down the guest operating system before it powers off the virtual machine to make the requested changes. The shut down process waits 120 seconds for a response from the target virtual machine, and fails without making changes to the virtual machine.

Solution

1 Check the target virtual machine in vCenter Server to determine if it has jobs running that are delaying the implementation of the action.

2 Retry the action from vRealize Operations Manager.

Troubleshooting Set CPU Count and Memory When Values Not Supported

If you run the Set CPU Count or Set Memory actions with an unsupported value on a virtual machine, the virtual machine might be left in an unusable state and require you to resolve the problem in vCenter Server.

Problem

You cannot power on a virtual machine after you successfully run the Set CPU Count or Set Memory actions. When you review the messages in Recent Tasks for the failed Power On VM action, you see messages stating that the host does not support the new CPU count or new memory value.

Cause

Because of the way that vCenter Server validates changes in the CPU and memory values, you can use the vRealize Operations Manager actions to change the value to an unsupported amount if you run the action when the virtual machine is powered off.

If the object was powered on, the task fails, but rolls back any value changes and powers the machine back on. If the object was powered off, the task succeeds, the value is changed in vCenter Server, but the target object is left in a state where you cannot power it on using the actions or in vCenter Server without manually changing the CPU or memory to a supported value.

Solution

1. In the menu, click **Administration**, then select **History** from the left pane.
2. Click **Recent Tasks**.
3. In the task list, locate your failed Power On VM action, and review the messages associated with the task.
4. Look for a message that indicates why the task failed.

   For example, if you ran a Set CPU Count action on a powered off virtual machine to increase the CPU count from 2 to 4, but 4 CPUs is not supported by the host. The Set CPU tasks reported that it completed successfully in recent tasks. However, when you attempt to power on the virtual machine, the tasks fails. In this example the message is `Virtual machine requires 4 CPUs to operate, but the host hardware only provides 2`.
5. Click the object name in the Recent Task list.

   The main pane updates to display the object details for the selected object.
6. Click the **Actions** menu on the toolbar and click **Open Virtual Machine in vSphere Client**.

   The vSphere Web Client opens with the virtual machine as the current object.
7. In the vSphere Web Client, click the **Manage** tab and click **VM Hardware**.
8. Click **Edit**.
9. In the Edit Settings dialog box, change the CPU count or memory to a supported value and click **OK**.

   You can now power on the virtual machine from the Web client or from vRealize Operations Manager.

Troubleshooting Set CPU Resources or Set Memory Resources When the Value is Not Supported

If you run the Set CPU Resources action with an unsupported value on a virtual machine, the task fails and an error appears in the Recent Task messages.

Problem

The Recent Tasks list indicates that a Set CPU Resource or Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see a message similar to the following examples.

```
RuntimeFault exception, message:[A specified parameter was not correct.
spec.cpuAllocation.reservation]
```

```
RuntimeFault exception, message:[A specified parameter was not correct. spec.cpuAllocation.limits]
```

Cause

You submitted the action to increase or decrease the CPU or memory reservation or limit value with an unsupported value. For example, if you supplied a negative integer other than -1, which sets the value to unlimited, vCenter Server could not make the change and the action failed.

Solution

- Run the action with a supported value.

  The supported values for reservation include 0 or a value greater than 0. The supported values for limit include -1, 0, or a value greater than 0.

Troubleshooting Set CPU Resources or Set Memory Resources When the Value is Too High

If you run the Set CPU Resources or Set Memory Resources action with a value that is greater than the value that your vCenter Server instance supports, the task fails and an error appears in the Recent Tasks messages.

Problem

The Recent Tasks list indicates that a Set CPU Resource or Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see messages similar to the following examples.

If you are working with Set CPU Resources, the information message is similar to the following example, where 1000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[1000000000] Mhz
```

The error message for this action is similar to this example.

```
RuntimeFault exception, message:[A specified parameter was not correct. reservation]
```

If you are working with Set Memory Resources, the information message is similar to the following example, where 1000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[1000000000] (MB)
```

The error message for this action is similar to this example.

```
RuntimeFault exception, message:[A specified parameter was not correct.
spec.memoryAllocation.reservation]
```

Cause

You submitted the action to change the CPU or memory reservation or limit value to a value greater than the value supported by vCenter Server, or the submitted reservation value is greater than the limit.

Solution

- Run the action using a lower value.

Troubleshooting Set Memory Resources When the Value is Not Evenly Divisible by 1024

If you run the Set Memory Resources action with a value that cannot convert from kilobytes to megabytes, the task fails and an error appears in the Recent Task messages.

Problem

The Recent Tasks list indicates that a Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see a message similar to the following example.

```
Parameter validation;[newLimitKB] failed conversion to (MB, (KB)[2000] not evenly divisible by 1024
```

Cause

Because vCenter Server manages memory reservations and limit values in megabytes, but vRealize Operations Manager calculates and reports on memory in kilobytes, you must provide a value in kilobytes that is directly convertible to megabytes. To do that, the value must be evenly divisible by 1024.

Solution

- Run the action where the reservation and limit values are configured with supported values.

  The supported values for reservation include 0 or a value greater than 0 that is evenly divisible by 1024. The supported values for a limit include -1, 0, or a value greater than 0 that is evenly divisible by 1024.

Troubleshooting Failed Shut Down VM Action Status

A shut down VM action task has a `Failed` status in the Recent Task list.

Problem

The Shut Down VM action did not run successfully.

The Recent Tasks list indicates that a Shut Down VM action has a task status of `Failed`. When you evaluate the Messages list for the selected job, you see `Failure: Shut down confirmation timeout`.

Cause

The shut down process involves shutting down the guest operating system and powering off the virtual machine. The wait time is 120 seconds to shut down the guest operating system. If the guest operating system does not shut down in this time, the action fails because the shut down action is not confirmed.

Solution

- Check the status of the guest operating system in vCenter Server to determine why is did not shut down in the allotted time.

Troubleshooting VMware Tools Not Running for a Shut Down VM Action Status

A Shut down VM action task has a `Failed` status in the Recent Task list and the Message indicates that VMware Tools were required.

Problem

The Shutdown VM action did not run successfully.

The Recent Tasks list indicates that a Shutdown VM action has a tasks status of `Failed`. When you evaluate the Messages list for the selected job, you see `VMware Tools: Not running (Not installed)`.

Cause

The Shutdown VM action requires that VMware Tools be installed and running on the target virtual machines. If you ran the action on more than one object, then VMware Tools was not installed, or installed but not running, on at least one of the virtual machines.

Solution

- In the vCenter Server instance that manages the virtual machine that failed to run the action, install and start VMware Tools on the affected virtual machines.

Troubleshooting Failed Delete Unused Snapshots Action Status

A Delete Unused Snapshots action task has a `Failed` status in the Recent Task list.

Problem

The Delete Unused Snapshots action did not run successfully.

The Recent Tasks list indicates that a Delete Unused Snapshots action has a tasks status of `Failed`. When you evaluate the Messages list for the selected job, you see this message.

```
Remove snapshot failed, response wait expired after:[120] seconds,
     unable to confirm removal
```

Cause

The delete snapshot process involves waiting for access to datastores. The wait time is 600 seconds to access the datastore and delete the snapshot. If the delete request is not passed to the datastore in that time, the action does not finish the delete snapshot action.

Solution

1. Check the status of the snapshot in vCenter Server to determine if it was deleted.
2. If it was not, submit the delete snapshot request at a different time.

Viewing Your Inventory

vRealize Operations Manager collects data from all the objects in your environment and displays a health, risk, and efficiency status for each object.

Survey your entire inventory to get a quick idea of the state of any object or click an object name for more detailed information. See Evaluating Object Information Using Badge Alerts and the Summary Tab.

Inventory Tab

The tab displays the state of each object in your environment. Objects are members of groups and applications that you define.

Where You Find Inventory

In the menu, click **Environment**, then select the **Inventory** tab.

Use the toolbar options to manage objects.

Table 5-77. Inventory Toolbar Options

Option	Description
Action	An action on the selected object. Depends 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 command 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.
Filter	Limit the list to objects matching the filter.

Table 5-78. Inventory Data Grid Options

Option	Description
Object Name	Displays a summary of the object.
Summary	Criticality of the health, risk, and efficiency of any object.

Planning the Capacity for Your Managed Environment Using vRealize Operations Manager

6

You can use the Projects feature in vRealize Operations Manager to plan for capacity allocations and upgrades in your virtual environment, or to optimize your existing resources. To plan your upcoming capacity needs, you create a project that anticipates forthcoming changes that affect the capacity of your objects.

In addition to creating projects to plan for hardware changes or virtual infrastructure changes, you can create custom profiles and custom data centers to help forecast your capacity needs. With custom profiles, you can determine how many instances of an object can fit in your environment depending on the available capacity and configuration. With custom data centers, you can see capacity analytics and badge computations based on the objects contained in the custom data center.

How Projects Work

A project is a detailed estimation of the capacity that you must have available in your environment based on upcoming changes. You can define projects to add or remove resources from objects such as your vCenter Server instance, clusters, data centers, hosts, virtual machines, and datastores.

With projects, you plan for changes in capacity, and examine the possible outcomes. You can plan for increases or decreases in the demand for capacity on your objects.

For example, if you plan to hire more staff in the next month, you must increase the capacity on the objects that they will use. To plan for this upcoming demand, you can create projects. In your projects, you add hosts to a data center, add memory and CPUs to a host, and increase the capacity of your virtual machines.

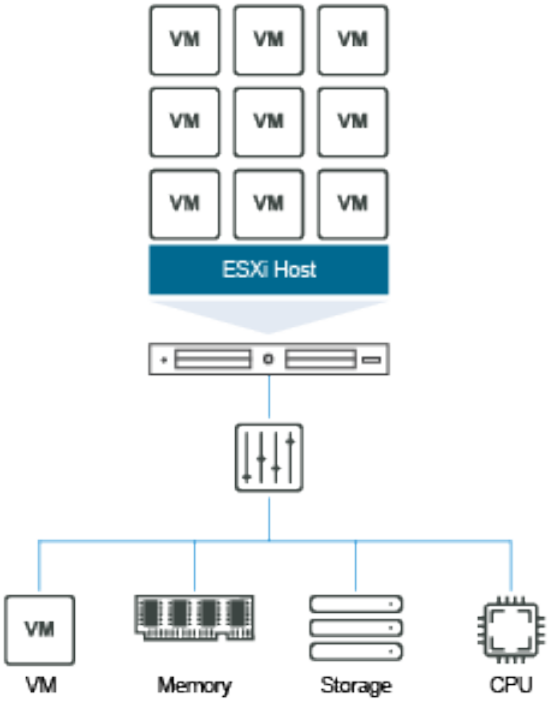

When you create a project, you add one or more capacity scenarios to the project to determine your future needs. Project scenarios anticipate the changes to capacity or demand that affect the object at an upcoming time and date. After you save each project, you drag the project to the visualization pane to chart the capacity forecast. You can see the anticipated capacity needs in the chart based on the values that you defined in your project scenarios. The visual representation shows how the needs for planned capacity compare to the resources that you currently have on those objects.

When you are sure that the objects require the planned capacity, you can commit the project to have vRealize Operations Manager reserve the capacity on those objects.

A project is a supposition about how the capacity and load change on your objects when you change the conditions in your virtual infrastructure environment. You do not have to implement the changes that your project represents. By creating the project, you can determine your capacity requirements before you implement the actual changes.

Projects List

The defined projects appear in a list below the visualization chart. vRealize Operations Manager filters the list according to the object that you select in the inventory tree. Use the toolbar to create, edit, or delete a project. To sort by columns in the list, click a column heading. To add a project to the visualization pane, click the plus icon, or drag the project to the pane between the list and the chart.

Visualization Chart

When you drag one or more projects into the visualization pane, the visualization chart displays each scenario that you defined in the projects.

The chart displays a numeric value for each scenario that you added to the project. For example, in a project for a host machine, the scenario named `Add Capacity: Percentage` is numbered 1.1, and the scenario named `Add Demand: Percentage` is 1.2.

To plan another host for your data center, you might also have a second project that includes a scenario named `Add Capacity: Add Host System`. The scenario in your second project is 2.1.

When you view both projects, the chart displays 1.1, 1.2, and 2.1 to indicate the point in time when each scenario takes effect.

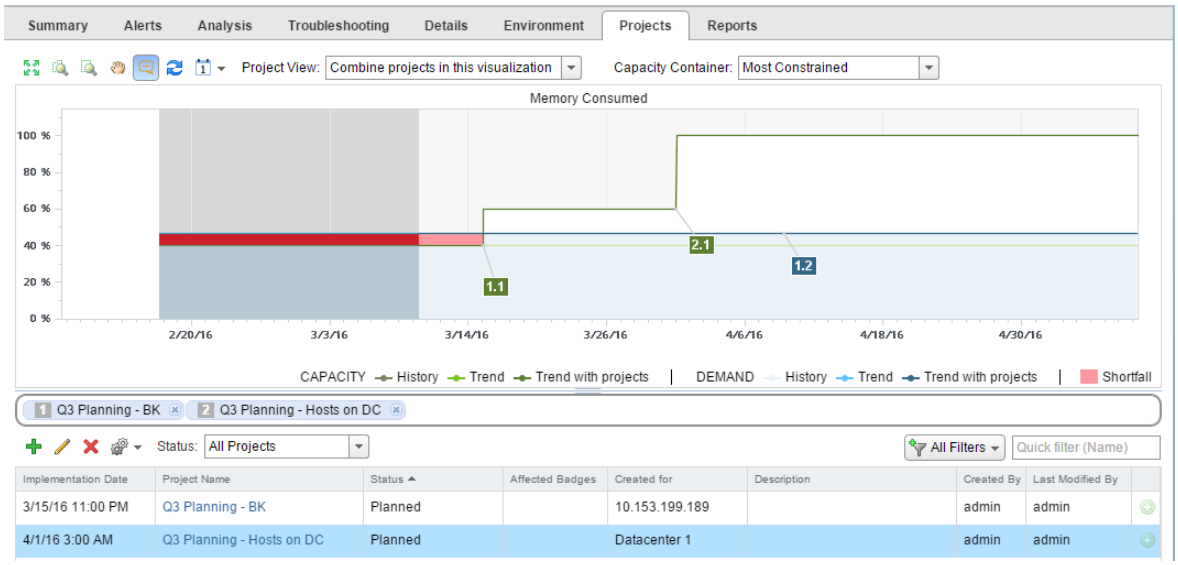

To view the details about the scenario, move the pointer to the number in the chart.

The projects and scenarios continue to appear in the chart until you delete them or refresh the view.

Project Scenarios Model Changes to Resources

You can use the following project scenarios to forecast capacity.

Table 6-1. Project Scenarios for Selected Objects

Selected Object	Project Scenarios
vCenter Server	Capacity ■ Add or remove host system, datastore, or percentage of capacity. ■ Change absolute capacity. Demand ■ Add or remove virtual machine or percentage of demand. ■ Change absolute demand.
Cluster	■ Add, remove, or update hosts. ■ Add, remove, or update datastores. ■ Add or remove virtual machines.
Host	Capacity ■ Add or remove datastore, or percentage of capacity. ■ Change absolute capacity. Demand ■ Add or remove virtual machine or percentage of demand. ■ Change absolute demand.
Datastore	Capacity ■ Add or remove percentage of capacity. ■ Change absolute capacity. Demand ■ Add or remove virtual machine or percentage of demand. ■ Change absolute demand.
Virtual Machine	■ Add, change, or remove capacity. ■ Add, change, or remove demand.

This chapter includes the following topics:

- Right-Sizing Capacity for Stress-Free Demand and Value
- User Scenario: Planning Capacity for an Increase in Workload
- Planning Hardware Projects in vRealize Operations Manager
- Planning Virtual Machine Projects and Scenarios
- Projects Tab in vRealize Operations Manager
- Custom Profiles in VMware vRealize Operations Manager
- Custom Datacenters in VMware vRealize Operations Manager

Right-Sizing Capacity for Stress-Free Demand and Value

Performance management and capacity planning vary across organizations and environments. Because the demand for capacity fluctuates in each environment, the top contenders for priority often include high efficiency versus low risk of poor performance. To plan and manage your

capacity needs and intelligently calculate the capacity of your resources, vRealize Operations Manager uses sophisticated models.

With the available capacity calculations, you can use various sophisticated models to produce practical correlations between objective measured metrics and subjective goals of acceptable performance and efficiency.

For example , the stress concept involves how high and how long the demand persists relative to the capacity available, and vRealize Operations Manager uses this value to measure the potential for performance problems. The higher the stress score, the worse the potential is for degraded performance on your objects. Depending on the configuration of the policy analysis settings for stress, a score of green might indicate 0–24 percent of stress. A score of red might indicate more than 50 percent of stress. With the five-minute data collections and the intelligent stress calculations, the system easily identifies periods of poor performance.

Demand drives stress. The system bases the calculations for right-sizing capacity on past demand. The goal of right-sizing is to produce a green stress level, marked by a green Stress badge.

Usable capacity is equal to the total capacity available minus any buffers that administrators or users defined. To measure the right-sized amounts of usable capacity, the capacity calculations use what is called a stress-free value. Using the demand, stress, and the stress-free value, vRealize Operations Manager calculates the right size.

The capacity analytics determine the actual and effective demand for resources based on having no contention. The calculations consider the capacity to be unlimited and free of contention for resources, which results in no stress on the available capacity. The result is called the stress-free demand or the stress-free value.

Where to Find Stress-Free Demand and Stress-Free Value

In some areas of the user interface, vRealize Operations Manager identifies capacity as Stress Free Demand, and in other areas it is identified as Stress Free Value. Both terms mean that the calculated capacity for an object is free from unacceptable levels of contention and stress, as defined in the policy for the Stress score.

Stress Free Demand appears in **All Metrics**, Views, and Reports.

- In **All Metrics**, you can use the metric named Stress Free Demand to examine the CPU demand, disk space allocation and demand, memory consumed, and the vSphere configuration limit on an object. When you apply this metric to these resources, you can build a metric graph to display the stress-free demand for an object. The graph displays the high and low stress-free capacity values over time.
- In **Reports**, you can use a view that includes the metric named Stress Free Demand to generate a report. The table in the report displays Stress Free Demand as the label. For example, this metric appears in the report named Cluster CPU Demand (%) Trend View.

Stress Free Value appears on the **Object > Analysis > Time Remaining** tab, and on the **Object > Analysis > Stress** tab.

- On the **Object > Analysis > Time Remaining** tab, you can view the time remaining for CPU demand, memory consumed, disk space demand and allocation, and the vSphere configuration limit. In this view, the table column name is Stress Free Value.
- On the **Object > Analysis > Stress** tab, the table column name is Stress Free Value. The tables display Stress Free Value as the calculated values for CPU demand, memory consumed, and the vSphere configuration limit.

Setting the Thresholds for the Stress Score

The analysis settings in the policy that you apply to your objects defines the thresholds for the stress score. The policy includes default settings for the stress score to be green, yellow, orange, or red. If the settings are too strict or loose for your environment, you can modify them.

To modify the stress score thresholds, edit the policy that applies to your objects, and click **Analysis Settings**. Select an object type and click the filter icon to display the policy analysis settings. In the Stress area, click the lock icon, expand **Stress**, and modify the stress thresholds.

In the analysis stress settings, vRealize Operations Manager uses the selected resources, such as Memory Demand, CPU Demand, and vSphere Configuration Limit to calculate the stress score.

You can set the stress thresholds to your own values, or turn them off. To change a stress score threshold, click and drag an icon along the slider. To remove a scoring range, such as the default range of 35–49 identified by orange, double-click an icon to disable the range.

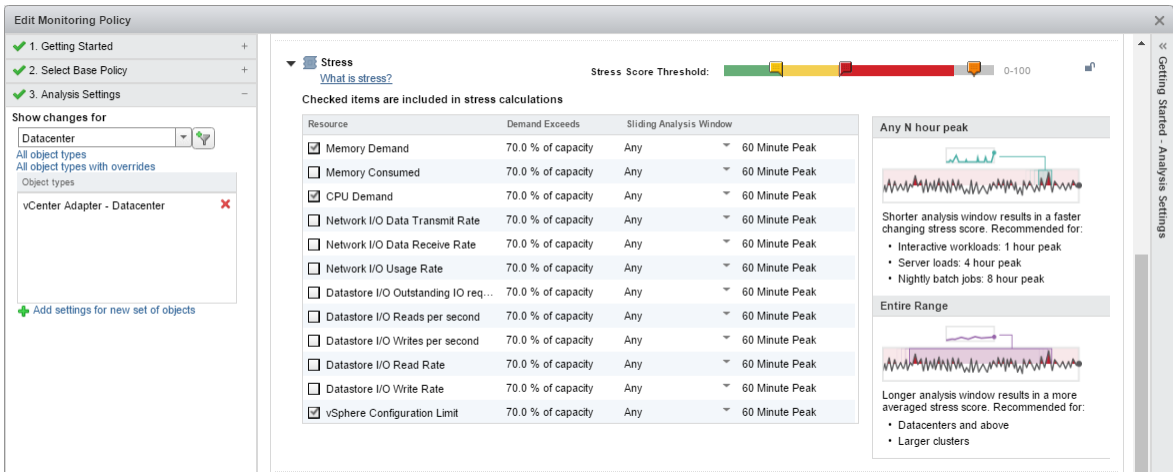

Demand Exceeds is a percentage of capacity. Capacity is also called provisioned capacity. To change the stress threshold for a resource, double-click the Demand Exceeds percentage, and enter the desired value. This value defines the point at which vRealize Operations Manager considers the percentage of demand to be stress. For example, to change the stress threshold for **Memory Demand**, double-click the current percentage, such as **70.0 % of capacity**, and enter the new percentage of demand to exceed for vRealize Operations Manager to identify stress.

For each resource, you can change the sliding analysis window value to include the entire range, and set the peak value to a different time depending on how you need vRealize Operations Manager to derive the stress score.

More About the Stress Score

vRealize Operations Manager calculates the stress zone and stress score for you. The following explanations cover typical scenarios where Demand does not exceed Capacity.

To determine the stress on an object for a specific time period, you can examine the demand curve to determine how much of the stress zone the demand occupies. The stress zone is typically where demand exceeds 70 percent of the total capacity. For example, stress occurs when CPU demand, memory demand, or memory consumed exceeds 70 percent of the capacity.

In a 60-minute peak period, vRealize Operations Manager bases the Stress score calculation on the following variables:

- Stress threshold, which is the `Demand Exceeds` setting
- Stress score threshold, which determines the color of the Stress badge
- Time range, as in 30 days of analysis
- Peak detection window, which is the 60-minute peak setting that you can adjust to either a non-zero number of minutes or the entire range.

When the demand exceeds 70 percent, that data point in time is in the Stress zone.

In the policy stress analysis settings, to examine an example graph used to calculate stress, click **What is stress?**.

Another example to explain the calculation used for CPU stress is shown here.

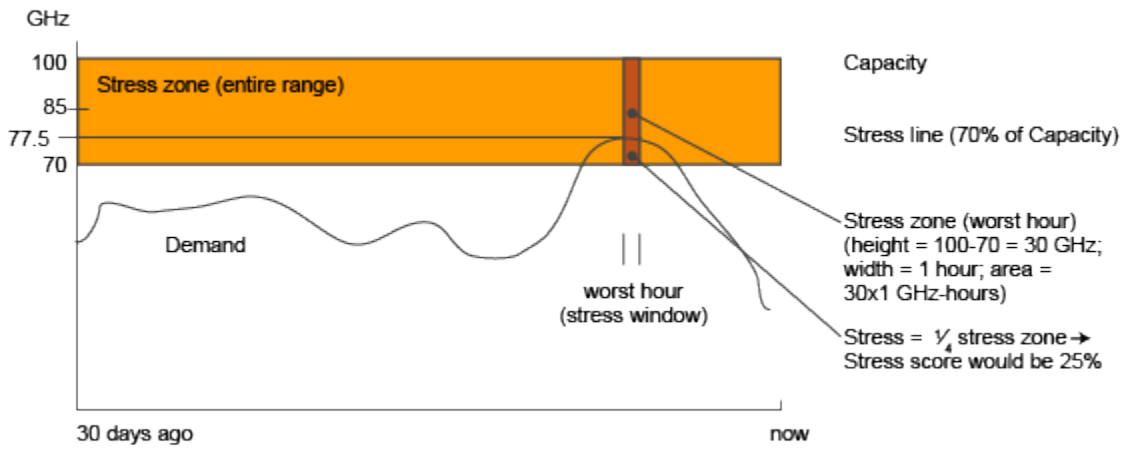

With a peak detection window size of 60 minutes, vRealize Operations Manager calculates the CPU stress score. It uses the area under the demand curve and above the stress threshold line as a percentage of the area covered by the total capacity curve.

Using time stamps of *t1* and *t2* to identify a 60-minute window in the last 30 days, the stress score depends on demand, stress threshold, and total capacity over time.

`Maximum((Demand – Stress Threshold) ÷ (Total Capacity – Stress Threshold))`

This equation applies to the stress calculations for each resource, such as memory demand, memory consumed, and CPU demand.

If `Total Capacity` varies during the time range being considered, `Stress Threshold` must also become variable, because `(Stress Threshold) = (Stress Threshold in %) × (Total Capacity)`.

Since `(Total Capacity)` can be a different value at a different time, as identified by *t*, then `“Stress Threshold”(t) = “Stress Threshold in %” × “Total Capacity”(t)`.

As a result, the Stress score is the highest aggregate of demand that exceeds 70 percent of capacity, as a percentage of the aggregate of capacity within any contiguous interval of 60 minutes in the last 30 days. The formula for the score is as follows:

`Maximum((Demand(t1, t2) – “Stress Threshold”(t1, t2)) ÷ (“Total Capacity”(t1, t2) – “Stress Threshold”(t1, t2)))`

Where:

- *t1* and *t2* are time stamps in the time continuum within the last 30 days.
- *t1* < *t2*
- *t2* - *t1* = 60 minutes
- Demand(*t1*, *t2*) is the demand curve between time *t1* and *t2*.
- “Stress Threshold”(*t1*, *t2*) is the stress threshold curve (as absolute values) between time *t1* and *t2*.
- “Total Capacity”(*t1*, *t2*) is the capacity threshold curve between time *t1* and *t2*.

vRealize Operations Manager calculates the aggregate during a contiguous time interval of 60 minutes in the last 30 days. The Stress score is the percentage of aggregate capacity in the same contiguous time interval of 60 minutes. An acceptable score yields a green Stress badge.

To view the Stress zone for an object, click **Object > Analysis > Stress**. Then, examine the Stress breakdown areas for CPU and memory, the Stress Zone column in the table, and the graph of actual demand.

By calculating the stress score, vRealize Operations Manager provides an intelligent way to evaluate peaks and fluctuations of the capacity of your objects over time.

User Scenario: Planning Capacity for an Increase in Workload

You are an IT administrator for one of your financial data centers. You must forecast the capacity requirements for your virtual infrastructure to plan for an increase in the workload of your cluster

and data center over the next month. To evaluate the demand and supply for capacity on your objects, and forecast the risk to your current capacity, you create projects and scenarios.

Your data center is named Fina_RDDC-01, and includes a cluster named Fina_RDCL-01. You plan to increase the overall workload on the cluster in this data center by 50 percent in the next month. You must also plan to add virtual machines and add one or more hosts to this cluster.

In this example, you create a project that includes scenarios to determine the impact of future capacity needs on your cluster objects. You then create a second project to plan for more capacity needs. Finally, you examine these projects together in the context of your current capacity so that you can understand the projected impact of these projects on your future capacity needs.

Prerequisites

Verify that vRealize Operations Manager has collected data for the last several weeks. See Connecting vRealize Operations Manager to Data Sources.

Procedure

1 Create a Sample Project to Increase Workload Capacity

   You are the IT administrator for the financial data center named Fina_RDDC-01 in your company. Create a project to plan for an increase in the workload on the cluster named Fina_RDCL-01 by 50 percent in the next month. In the project, you create scenarios that anticipate the effect of the capacity needs on the hosts, virtual machines, and cluster in the data center.

2 Create a Sample Project to Add a Host and Virtual Machines

   You are the IT administrator for the financial data center in your company. To plan for capacity needs on the cluster named Fina_RDCL-01 in the data center named Fina_RDDC-01, you create another project. In your project, you add virtual machines and a host to the cluster.

3 View the Result of Your Capacity Projects

   You are the IT administrator responsible for the data center named Fina_RDDC-01. You view the effects of the projects and scenarios that you created on the overall capacity of the cluster in your data center.

Create a Sample Project to Increase Workload Capacity

You are the IT administrator for the financial data center named Fina_RDDC-01 in your company. Create a project to plan for an increase in the workload on the cluster named Fina_RDCL-01 by 50 percent in the next month. In the project, you create scenarios that anticipate the effect of the capacity needs on the hosts, virtual machines, and cluster in the data center.

Use your new project and scenario to determine what happens to the capacity of the objects in your environment when you plan for an increase in demand.

Prerequisites

- Understand the scope of this sample workflow. See User Scenario: Planning Capacity for an Increase in Workload.
- Verify that the cluster named `Fina_RDCL-01` in your data center named `Fina_RDDC-01` includes multiple hosts and virtual machines.

Procedure

1. In the menu, click Environment, then click Custom Data Centers.
2. In custom data center the inventory tree, select the data center named `Fina_RDDC-01`. Then select the cluster named `Fina_RDCL-01`.
3. Click the **Projects** tab.
4. On the toolbar above the Projects list pane, click **Add**.
5. In the Projects workspace, enter a name and description for the project.

   For example, `Fina RDCL Q1 Planning`.
6. For the Status, select **Planned - no badges affected**.
7. In the workspace, click **Scenarios**.
8. Under Add Demand, drag the scenario named **add percentage of demand** to the Scenarios pane.

   The scenario is numbered 1.1.
9. In the Configuration pane, configure the demand.
   a. Click the **Implementation Date** calendar icon, and select the date one month from today.
   b. In the Use Global Value text box, enter `50`.
10. To add the scenario to your project, click **Save** and click **Close**.

Results

vRealize Operations Manager saves the scenario in the project.

What to do next

To add virtual machines and hosts to the cluster named `Fina_RDCL-01`, create another project and scenario. See Create a Sample Project to Add a Host and Virtual Machines.

Create a Sample Project to Add a Host and Virtual Machines

You are the IT administrator for the financial data center in your company. To plan for capacity needs on the cluster named `Fina_RDCL-01` in the data center named `Fina_RDDC-01`, you create another project. In your project, you add virtual machines and a host to the cluster.

You create another project to add a host and virtual machine to the cluster named `Fina_RDCL-01` so that you can see the effect on the capacity of the cluster. The cluster already includes several hosts named `Fina_RDH-01` and `Fina_RDH-02`.

Prerequisites

Create a project to plan for an increase in the workload on the cluster named `Fina_RDCL-01` by 50 percent in the next month. See Create a Sample Project to Increase Workload Capacity.

Procedure

1. In the menu click Environment, then click Custom Data Centers.
2. In the Custom Data Centers inventory tree, select the data center named `Fina_RDDC-01`, and the cluster named `Fina_RDCL-01`
3. Click the **Projects** tab.
4. On the toolbar above the Projects list pane, click **Add**.
5. In the Projects workspace, enter a name and description for the project.

   For example, `Fina RDCL-01 Hosts_VMs Q1 Planning`.
6. For the Status, select **Planned - no badges affected**.
7. In the workspace, click **Scenarios**.
8. Under Add Demand, drag the scenario named **add Virtual Machine** to the Scenarios pane.

   The scenario is numbered 1.1.
9. In the Configuration pane, configure the capacity requirements.
   a. Under Changes, enter **10** for the number of virtual machines.
   b. Under Metrics, enter **4 GB** for Memory (Consumed).
   c. For CPU - Allocation model for vCPUs, enter **2**.
10. Under Add Capacity, drag the scenario named **add Host System** to the Scenarios pane.

    The scenario is numbered 1.2.
11. In the Configuration pane, configure the host.
    a. Under Changes, enter **2** for the number of hosts.
    b. Under Metrics, enter **8 GB** for Memory Demand.
    c. For CPU Allocation, enter **4** for the number of vCPUs.
12. To add the scenario to your project, click **Save** and click **Close**.

Results

vRealize Operations Manager saves the scenario in the project.

What to do next

Visualize the effect of your capacity planning projects in the visualization chart. View the Result of Your Capacity Projects.

View the Result of Your Capacity Projects

You are the IT administrator responsible for the data center named `Fina_RDDC-01`. You view the effects of the projects and scenarios that you created on the overall capacity of the cluster in your data center.

View both of your projects so that you can visualize the anticipated requirements simultaneously. Use the results to plan your overall capacity needs for the cluster named `Fina_RDCL-01` in the data center named `Fina_RDDC-01`.

Prerequisites

Create a project so that you can plan to add hosts and virtual machines to the cluster named `Fina_RDCL-01`. See Create a Sample Project to Add a Host and Virtual Machines.

Procedure

1. Select your cluster named `Fina_RDCL-01`, and click the **Projects** tab.
2. In the Projects list, select the project named `Fina RDCL Q1 Planning`, and drag it to the pane just above the Projects list.
3. Select the project named `Fina RDCL-01 Hosts_VMs Q1 Planning`, and drag it to the pane just above the Projects list.
4. To view both projects in the visualization chart, from the Project View drop-down menu above the chart, select **Combine projects in this visualization**.

Results

The combined values for your projects appear in the visualization chart.

What to do next

Determine whether to commit the projects so that you can reserve the capacity on the objects in your data center.

Planning Hardware Projects in vRealize Operations Manager

Planning a capacity project for the hardware in your infrastructure involves changes to the host hardware and datastore hardware. To determine whether you must purchase new hardware, you can create projects.

Before you change your hardware objects, you can create and implement a hardware project to determine the result of the change. With hardware projects, you can determine the capacity requirements for your objects before you change the hardware in your environment.

You might need to plan for hardware changes under various circumstances.

- If you implement new applications, you must ensure that your objects have enough resources to support the amount of disk space required after you deploy the applications.
- If you add hosts to an existing cluster, you must ensure that the cluster can sustain the increase in capacity used during the following quarter of the year.
- If you make a configuration change to the demand for memory or CPU on your objects, you must understand the capacity requirements and workloads of your existing objects.

Create a Project to Plan for Hardware Changes

To support an increase in the capacity requirements for the objects in your environment, you can create projects to determine whether a purchase of new hardware is necessary.

To forecast the capacity requirements for your objects when you add, update, or remove hardware capacity, you create projects and add scenarios to those projects. This procedure creates a hardware project that forecasts changes to a host in your cluster.

Prerequisites

vRealize Operations Manager has collected data for the last several weeks. See Connecting vRealize Operations Manager to Data Sources.

Procedure

1. In the menu click Environment, then click Inventory and select a host in the tree.

   Alternatively, in the left pane drill down to find the object you want.
2. Click the **Projects** tab.
3. On the toolbar above the visualization area, from the Capacity Container drop-down menu, click **Most Constrained**.
4. On the toolbar below the visualization area, click **Add**.
5. In the Projects workspace, enter a name and description for the project.
6. For the Status, select **Planned - no badges affected**.
7. In the workspace, click **Scenarios**.
8. Under Add Capacity, drag the scenario named **add Datastore** to the Scenarios area.
9. In the Configuration area, enter the general parameters for the project scenario.

Option	Description
Implementation Date	Set the date and time to implement the project scenario.
Changes	Set the number of datastores to add.

Option	Description
Populate metrics from	Copy the disk space use and allocation metrics from an existing datastore, and select an existing datastore.
Metrics	Set the amount of disk space use and allocation.

10 To view the effect of your selections in the visualization chart, click **Save project and continue editing**.

With the Capacity Container set to **Most Constrained**, the visualization chart might indicate a CPU shortfall when you implement the project scenario. This shortfall might occur because the CPU allocation might be greater than the available capacity. In this case, you might need to add CPU capacity before you implement the project scenario.

11 When you are satisfied with the capacity forecast based on your settings, click **Save** to add the scenario to the project.

12 On the Projects tab, click your project in the list and drag it to the area above the project list.

Results

vRealize Operations Manager applies your project and scenario to the visualization chart. The capacity forecasted in the project appears as a gray line in the chart.

What to do next

Add the scenario named **Add Demand: add percentage of demand** to the project, and set the Capacity Container to **Disk Space Allocation**. The visualization chart might indicate that when you implement the project scenario, you have a disk space shortfall. In this case, you might need to add disk space capacity before you implement the project scenario.

In the visualization chart, evaluate the current available capacity with the actual capacity required if you change your environment as defined in your project. Determine whether to commit the project so that it reserves the capacity required for the hardware change.

Planning Virtual Machine Projects and Scenarios

Virtual machine projects help you assess the consequences of changing resources on virtual machines without applying the changes to your virtual environment. Before you apply changes to your virtual environment, you can create sample virtual machine projects to model adding or removing virtual machines to a host or a cluster.

- Create a Virtual Machine Project Using Populated Metrics

  You can create a project scenario that uses an existing virtual machine profile as a model. The project scenario simulates the resource requirements when you add one or more virtual machines to a host or cluster.

- Create a Sample Project for a New Virtual Machine

  Virtual machine projects assess the consequences of adding a new virtual machine to a cluster or host, without applying the actual changes to your virtual environment.

- Create a Sample Project to Simulate Removing a Virtual Machine

  You can create a project that simulates removing one or more virtual machines from a host or a cluster. You might remove virtual machines when you no longer need them, or when you must move them.

Create a Virtual Machine Project Using Populated Metrics

You can create a project scenario that uses an existing virtual machine profile as a model. The project scenario simulates the resource requirements when you add one or more virtual machines to a host or cluster.

When you configure the settings in a project scenario to add virtual machines, you can populate the resource values for the planned virtual machine from an existing profile. Or, you can copy the values from an existing virtual machine.

To calculate the capacity metrics values for the virtual machine, vRealize Operations Manager partitions the capacity for CPU, memory, and disk dimensions, according to the profile that you select.

For information about CPU and memory maximums, see the VMware vSphere documentation.

Procedure

1. In the menu click **Environment**, then click **Inventory**.

   Alternatively, in the left pane drill down to find the object you want.

2. Click the host or cluster that contains the planned virtual machine reside.
3. Click **Projects**.
4. Click **Add New Project**.
5. In the Projects workspace, enter a name and description for the project.
6. For the Status, select **Planned - no badges affected**.
7. In the workspace, click **Scenarios**.
8. Under Add Demand, drag the scenario named **add Virtual Machine** to the Scenarios area.

9 In the Configuration area, enter the general parameters for the project scenario.

a Select the date and time to implement the project scenario.

b Click **Populate metrics from**, select an existing profile or an existing virtual machine, and click **OK**.

Option	Action
Copy metric values from a pre-defined profile.	From the Profile drop-down menu, select an existing profile to populate the metrics values for the planned virtual machine.
Copy metric values from an existing object.	From the Existing Virtual Machine drop-down menu, select a virtual machine to populate the metrics values for the planned virtual machine. The list displays the virtual machines that reside on the selected object.

c (Optional) To duplicate virtual machines, increase the virtual machine count.

d To see the effect of the planned virtual machines in the visualization chart, click **Save project and continue editing**.

With the Capacity Container set to **Most Constrained**, the visualization chart might indicate that you have a CPU shortfall when you implement the project scenario. The shortfall might occur because the CPU allocation might be greater than the available capacity. In this case, you might need to add CPU capacity before you implement the project scenario.

10 When you are satisfied with the capacity forecast based on your settings, click **Save** to add the scenario to the project.

11 On the Projects tab, click your project in the list and drag it to the area above the project list.

Results

vRealize Operations Manager applies your project and scenario to the visualization chart. The capacity forecasted in the project appears as a gray line in the chart.

What to do next

In the visualization chart, evaluate the current available capacity with the actual capacity required if you change your environment as defined in your project. Determine whether to commit the project so that it reserves the capacity required for the new virtual machines.

Create a Sample Project for a New Virtual Machine

Virtual machine projects assess the consequences of adding a new virtual machine to a cluster or host, without applying the actual changes to your virtual environment.

For information about relevant CPU and memory maximums, see the VMware vSphere documentation.

Procedure

1 In the menu click Environment, then click Inventory and select a destination object in the tree.

   Alternatively, in the left pane drill down to find the object you want.

   If you implement your scenario, the destination object is a cluster or host where you locate the new virtual machines.

2 Click the **Projects** tab and click the **Add New Project** icon.

3 From the Projects workspace, enter the name and a description of the project.

4 Select the **Planned** status.

5 To add scenarios to this project, click **Scenarios**.

6 Select the **add Virtual Machine** scenario and drag it to the Scenarios area.

7 Set the virtual machine count and the configuration for the virtual machine.

   vRealize Operations Manager does not require you to set the disk I/O and network I/O use of the new virtual machines. vRealize Operations Manager uses the average disk I/O and network I/O use across virtual machines in the host or cluster as an estimation of the new virtual machine use.

8 To see the effect in the visualization chart when your configuration selections are finished, click **Save project and continue editing**.

9 To add the scenario to the project, click **Save**.

10 To close the Project workspace, click **Close**.

   Clicking **Close** discards all changes. Clicking **Save project and continue editing** persists any changes that were not previously saved.

Results

vRealize Operations Manager applies the project to the object you selected. The project shows the current capacity compared to the expected capacity when you add the virtual machines to the target object.

Create a Sample Project to Simulate Removing a Virtual Machine

You can create a project that simulates removing one or more virtual machines from a host or a cluster. You might remove virtual machines when you no longer need them, or when you must move them.

Procedure

1 In the menu, click Environment, then click Inventory and select a host or cluster from the tree.

2 Click the **Projects** tab.

3 On the toolbar below the visualization area, click **Add**.

4 In the Projects workspace, enter a name and description for the project.

5 For the Status, select **Planned - no badges affected**.

6 In the workspace, click **Scenarios**.

7 Under Remove Demand, drag the scenario named **remove selected object** to the Scenarios area.

8 In the Configuration area, under Changes, click **Select one or more objects to remove**.

9 From the list of objects, click the check box for a **Virtual machine**, and click **OK**.

10 To add the scenario to the project, click **Save**.

11 On the Projects tab, click your project in the list and drag it to the area above the project list.

Results

vRealize Operations Manager applies your project and scenario to the visualization chart. The capacity forecasted in the project appears as a gray line in the chart. Compare the current capacity to the expected capacity if you commit this project to remove one or more virtual machines from the selected object.

What to do next

You can create other projects, and combine or compare the outcomes in the visualization chart.

Projects Tab in vRealize Operations Manager

The **Projects** tab is a list of all the projects generated for the selected object, group, or application. You can create projects, access existing projects, and view the capacity trend of historical data on the project visualization chart.

How the Projects Tab Works

On the **Projects** tab, you create projects and add scenarios to those projects so that you can forecast the capacity of your objects. The objects can include vCenter Server instances, clusters, hosts, datastores, and virtual machines. In the visualization area, when you add or remove projects, vRealize Operations Manager displays the cumulative effect of those projects on the object selected in the inventory tree.

Where You Find the Projects Tab

To create and modify projects:

- In the menu, click **Environment**, then select a group, custom data center, application or inventory object. Click the **Projects** tab.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to quickly drill down to the object you want.

Table 6-2. Projects Tab

Options	Description
Projects visualization area and toolbar	Use the Project View drop-down menu to select the way vRealize Operations Manager displays the projects. In the visualization chart, the project views assign names to the projects and scenarios, such as 1.1, 1.2, and 2.1. ■ **Combine projects in this visualization**. Combines projects in one graph. ■ **Compare projects in this visualization**. Shows each project in a separate smaller graph. Use the Capacity Container drop-down menu to select a container for this project. The container options change depending on the object you select. For example, for a cluster, you can forecast capacity according to most constrained, memory or CPU demand, a vSphere configuration limit, disk space allocation, or disk space demand. The visualization chart, which displays stress-free demand and usable capacity, includes other metrics in addition to the metrics that you modify. As a result, the magnitude of the change in capacity might not scale proportionally to your input. Use the visualization area toolbar options to zoom and pan the view in various ways, show data values, refresh the chart, and display the range of data.
Projects list toolbar	Use the projects pane toolbar selection to manage your projects. You can add a project, edit the configuration of an existing project, and remove a project from the list. To change the status of a selected project, click the gear. ■ **Change status to Planned**. Sets the created project to the planned state, and runs the what-if analysis to visualize the forecasted effects on the capacity of your object. ■ **Change status to Committed**. Commits or reserves the project in the selected capacity container. All views, reports, and dashboards reflect the project capacity as though you deployed the project. To filter the projects list, click the **Status** drop-down menu. You can filter the list of projects, and sort the columns in the data grid.
Projects list	The object that you select in the inventory tree determines the projects that populate in the projects list. All projects that appear are associated directly with the selected object or with its children. To add a project to the visualization area, drag the project row to the area above the Projects list, or click the plus icon in the project row.

Projects Name and Description Workspace

You use the Projects workspace to create projects, which represent upcoming environment changes that affect the capacity of the object. You define the project name, add a description, and select a status. You add one or more scenarios to the project to forecast the change to capacity that you expect to implement.

Where You Define the Project

To create, edit, view, and forecast a project, click **Environment** in the menu, select an object, and click the **Projects** tab. On the Projects toolbar, click the plus sign to add a project. To edit a selected project, click the edit icon.

Options	Description
Name	Name of the project, which appears on the **Projects** tab.
Description	Meaningful description of the project.
Status	■ **Planned - no badges affected**. Sets the created project to the planned state, and runs the what-if analysis to visualize the forecasted effects on the capacity of your object. ■ **Committed - badges affected**. Commits or reserves the project in the selected capacity container. All views, reports, and dashboards reflect the project capacity as though you deployed the project. To determine whether the reserved resources for this project affect the time remaining or both the time and capacity remaining, click **Advanced** and select one of the menu items. When you select **Committed - badges affected**, capacity is reserved based on the advanced setting that you select. In the advanced settings, when you select **This project affects the Time Remaining badge**, vRealize Operations Manager reserves capacity on the implementation date set for the project. When you select **This project affects the Time Remaining badge**, vRealize Operations Manager reserves capacity immediately.

Projects Scenarios Workspace

A project scenario is a simulation about how capacity changes when you change the conditions to forecast the upcoming capacity to your virtual infrastructure. Project scenarios do not make the actual changes to objects in your environment. But when you implement the scenario, you can determine the capacity requirements before you must change your environment.

Where You Add Scenarios

To add a scenario to a project, or to update a scenario in an existing project, click **Environment** in the menu, select an object, and click the **Projects** tab. On the Projects toolbar, click the plus sign to add a project, or click an existing project and click the pencil to edit the project. In the Projects workspace, click **Scenarios**.

Options	Description
Object	From the drop-down menu, double-click an object to select it or use the filter to find an object. The selected object determines the content of the projects scenario list.
List of scenarios	To add the scenario to the project, drag a scenario to the Scenarios area. When you add, edit, or remove a project scenario, to view your changes in the visualization chart click **Save project and continue editing**.
Capacity Container	From the drop-down menu, select a container for this scenario. Use the Capacity Container drop-down menu to select a container for this project. The container options change depending on the object you select. For example, for a cluster, you can forecast capacity according to most constrained, memory or CPU demand, a vSphere configuration limit, disk space allocation, or disk space demand.

Options	Description
Visualization Chart	The visualization chart, which displays stress-free demand and usable capacity, includes other metrics in addition to the metrics that you modify. As a result, the magnitude of the change in capacity might not scale proportionally to your input. The what-if visualization chart uses average hourly data. The Current Values settings in the project configuration reflect the latest 5-minute data points. For example, the Current Values settings appear in the scenarios named `change absolute capacity` and `change absolute demand`. For a metric, a large variance can occur between the average hourly data and the latest data point. When you change a value according to the latest data point, the visualization chart displays the change according to the hourly average. For example, the memory demand hourly average is 35 GB, but the latest data point in Current Values in the project configuration dropped to 3.5 GB. You can use the Change by Absolute scenario to change the demand to 7 GB, with the intent to double the demand. In this case, the visualization chart renders this change as a drop in demand from its average of 35 GB.
Configuration	Configure the following information. ■ Scenario Name ■ Scenario Description. A meaningful description of the scenario. ■ Implementation Date. The date and time to implement the scenario. ■ Configuration area. Depending on the scenario, configure the global value, or customize the metrics and metric values. The Changes area varies depending on the scenario that you configure. If you add a scenario to a virtual machine or host object, you can populate the metrics from an existing object or from a custom profile. For example, to apply to the capacity calculations when you configure the `Add Demand: add Virtual Machine` scenario, you can click **Populate metrics from**. You copy the metric values from a predefined profile or from an existing object. When you copy the metric values from an existing object, you can either use the most recent metric values or a historical demand pattern from another virtual machine.

Custom Profiles in VMware vRealize Operations Manager

A custom profile is a user-defined instance of the capacity allocation and demand for a specific object type. You can use custom profiles to help forecast the capacity needs for your environment.

To determine how many instances of the object can fit in your environment, use custom profiles with projects and scenarios. Depending on the available capacity in your environment, you can add one or more instances of the object that the custom profile capacity requirements represent.

When you create a custom profile for an object type, such as a virtual machine, you create a project and add a virtual machine scenario to it. In the project scenario, you select your custom profile to populate the metrics and capacity for that object type to the project scenario. You use the capacity sizing of your custom profile to forecast the capacity needs for the parent object of the virtual machine.

To determine how many instances of the custom profile object you can include on the parent object, you select the parent object, click **Analysis**, and click **Capacity Remaining**. The custom profiles appear on the What Will Fit section of the Capacity Remaining Breakdown area, and indicate how many instances of the object fit in your environment.

Custom Profiles Details and Related Policies

A custom profile defines a specific configuration of an object instance. With profiles, you can determine how many instances of that object can fit in your environment, depending on the capacity available and the configuration of that object instance.

How Custom Profiles Work

As with default profiles, custom profiles define metrics configurations for an object. You can create as many custom profiles as you need for an object type. For example, you might create one custom profile for a virtual machine that has a memory demand model of 2 GB. You create another custom profile that has a memory demand model of 4 GB.

vRealize Operations Manager uses custom profiles of virtual machines to calculate the number of virtual machines that can fit in your environment. The number of virtual machines is based on the capacity allocation and demand defined in the profile. To examine the capacity calculations, select a parent object such as a host or cluster. Click **Analysis > Capacity Remaining**, and view the What Will Fit section of the Capacity Remaining Breakdown area.

You can also use custom profiles to populate metrics when you create project scenarios. To use a custom profile in a project scenario, select an object such as a host or cluster. Click **Projects**, and click **Add** to create a project. When you add a scenario to your project, such as to add a virtual machine, you click **Populate metrics from**. You select your custom profile to include the capacity settings defined in your custom profile to the project scenario. See also Projects Scenarios Workspace.

Where You Find Custom Profiles

In the menu, click **Administration** then **Configuration > Custom Profiles** in the left pane.

Table 6-3. Custom Profiles Options

Option	Description
Toolbar options	Use the toolbar options to manage custom profiles. ■ **Add New Profile**. Add a custom profile for a specific object type. ■ **Edit Selected Profile**. Modify the selected profile. ■ **Delete Selected Profile**. Remove the selected profile. ■ **Clone Selected Profile**. Create a copy of the selected profile and customize it for your needs.
Filtering options	Filter the list to display profiles that match the filter you create. You can sort by name, description, object type, or adapter type. Or, enter filter text in the Quick filter text box.

Table 6-3. Custom Profiles Options (continued)

Option	Description
Details tab	Displays the name, description, adapter, object type, and metrics applied to the custom profile.
Related Policies tab	Displays all the policies associated with the selected custom profile. To modify the policies associated with the custom profile, edit the profile. If **Enable this profile for all Policies** is selected, deselect it, and click the **x** to remove the policies not to be associated with the custom profile.

Custom Profiles Add and Edit Workspace

You can add a custom profile for an object type to determine how many instances of a specific object can fit in your environment. In the Custom Profiles workspace, you create a custom profile for an object and define its capacity configuration.

Where You Create or Edit a Custom Profile

To create a custom profile, click **Administration** in the menu, then **Configuration > Custom Profiles** in the left pane. To create a custom profile, click the plus sign. To edit the selected profile, click the edit icon. To use an existing profile as a template, click **Clone Selected Profile**.

Table 6-4. Custom Profiles Configuration Options

Option	Description
Profile Name	Descriptive name of the custom profile.
Profile Description	Meaningful description for the custom profile. Provide specific information that other users must know about this profile.
Object Type	Basic object for the profile, such as a virtual machine.
Enable this profile for all Policies	Used to override all other policy settings. To display the list of available policies, and select individual policies from the list, deselect this option.
Advanced	Displays the policy and blacklist menu items. ■ Enable for Policy. Lists the policies enabled for use with the custom profile. You can remove policies from the list, and select only the policies to use with the custom profile. ■ Hide Profile from. Displays objects for which the custom profile does not apply. To add more than one object type from which to hide the custom profile, click **Add Blacklist Object Type**, and select the object type from the list.
Metrics	Capacity requirements for the object instance, based on the metrics you specify. You can use an existing object or profile to populate the capacity metrics.
Filter (Model)	Filters the capacity metrics by allocation or demand to determine the capacity available or required by the object. For example, you can view only the CPU and memory allocation, or their demand, or both. The default model is allocation.

Custom Datacenters in VMware vRealize Operations Manager

A custom data center is a user-defined container for a group of objects that includes clusters, hosts, and virtual machines. Custom data centers provide capacity analytics and capacity badge computations based on the objects it contains. You can use custom data centers to forecast and analyze the capacity needs for your environment.

When you create a custom data center, you can include multiple cluster objects that span multiple vCenter Server instances. For example, you might have a production environment that spans multiple clusters, and you must monitor and manage the performance and capacity of the entire production environment.

After you create your custom data center, you can select it in the list of custom data centers to display a summary of its health, risk, and efficiency. To access the list of custom data centers, click **Environment** on the top menu.

This view displays the top alerts for the data center. To examine the capacity remaining for the custom data center, click the **Analysis** tab, and click **Capacity Remaining**.

You can use your custom data center objects to balance the workload across the clusters in your environment. Click **Home**, click **Dashboard List**, click the dashboard named **Workload Distribution**, and view the use of your custom data center in the dashboard.

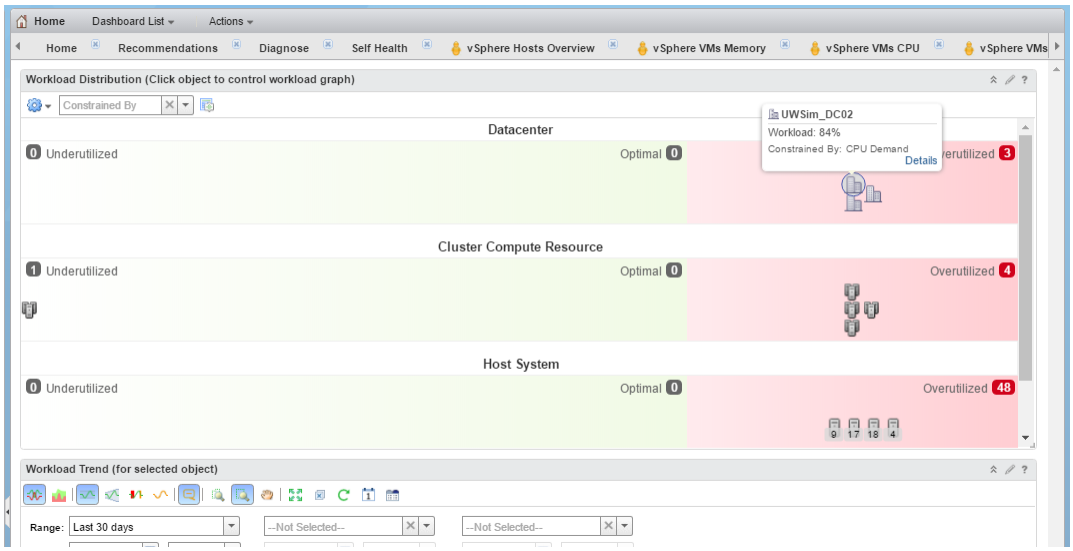

Click the icon for your data center to view its workload trend, CPU and memory workload measurements, and the vSphere configuration limit.

Custom Datacenters List

You can view the list of custom data centers that exist in your environment, and a summary view of its health, risk, and efficiency. In this view, you can click a custom data center to display the top alerts that the objects in the custom data center triggered.

How Custom Datacenters Work

In vSphere, a data center serves as a container for objects that a vCenter Server instance manages. A custom data center is a container that can include objects from multiple vCenter Server instances.

Custom data centers can contain vCenter Server instances, data centers, clusters, hosts, virtual machines, and datastores. You can add vSphere object types to a custom data center.

When you add an object, the hierarchical children of that object become part of the custom data center. An object can belong to multiple custom data centers.

When you create custom data centers, the system runs capacity analytics on the objects in the custom data center, even if those objects span multiple vCenter Server instances. For example, you might need to examine the capacity analytics data across multiple clusters, and the multiple vCenter Server instances that manage those clusters. You do not have to analyze the capacity of one cluster or one vCenter Server instance at a time. You can create a custom data center, add all the clusters to it, and see the capacity analysis in a single location.

Where You Find Custom Datacenters

Select **Environment** in the menu and click the **Custom Datacenters** tab.

Table 6-5. Custom Datacenters Toolbar and Grid Options

Option	Description
Toolbar options	Use the toolbar options to manage your custom data centers. ■ **Add New Custom Datacenter**. Add a custom data center. ■ **Edit Custom Datacenter**. Modify the selected custom data center. ■ **Delete Custom Datacenter**. Remove the selected custom data center. ■ **Clone Custom Datacenter**. Create a copy of the selected custom data center and customize it for your needs.
Filter	Limit the list of custom data centers to those data centers that match the text that you enter in the **Filter** text box.
Data grid	Lists the custom data centers in your environment, and displays the health, risk, and efficiency for each one. To view a summary of the custom data center health, risk, and efficiency on the **Summary** tab, click the custom data center name. To edit, delete, or clone a custom data center, click to the right of the custom data center name. Then, click the toolbar option.

Custom Datacenters Add and Edit Workspace

A custom data center is an object type that provides capacity analytics and capacity badge computations based on the objects it contains. You create a custom datacenter object and add inventory objects to it.

Where You Create or Edit a Custom Datacenter

To create a custom data center, in the menu click **Environment**, click the **Custom Datacenters** tab, and click the plus sign.

To edit a selected custom data center, click to the right of the custom data center name, and click the edit icon. To use an existing custom data center as a template, click to the right of the custom data center name, and click the clone icon.

Table 6-6. Add and Edit Custom Datacenters Configuration Options

Option	Description
Name	Descriptive name of the custom data center.
Description	Meaningful description for the custom data center. Provide specific information that other users must know about this custom data center.
Objects	Lists the objects in your environment. Select the check box for each object to add to the custom data center. You can add vCenter Server instances, vSphere data centers, vSphere clusters, and ESXi hosts. When you add an object, the hierarchical children of that object become part of the custom data center. An object can belong to multiple custom data centers.

Metric, Property, and Alert Definitions

7

vRealize Operations Manager provides definitions for the metrics, properties, and alerts defined on objects in your environment.

This chapter includes the following topics:

- Metric Definitions in vRealize Operations Manager
- Alert Definitions in vRealize Operations Manager
- Property Definitions in vRealize Operations Manager

Metric Definitions in vRealize Operations Manager

Metric definitions provide an overview of how the metric value is calculated or derived. If you understand the metric, you can better tune vRealize Operations Manager to display results that help you to manage your environment.

vRealize Operations Manager collects data from objects in your environment. Each piece of data collected is called a metric observation or value. vRealize Operations Manager uses the VMware vCenter adapter to collect raw metrics. vRealize Operations Manager uses the vRealize Operations Manager adapter to collect self-monitoring metrics. In addition to the metrics it collects, vRealize Operations Manager calculates capacity metrics, badge metrics, and metrics to monitor the health of your system.

All metric definitions are provided. The metrics reported on your system depend on the objects in your environment. You can use metrics to help troubleshoot problems. See Troubleshooting with the All Metrics Tab.

Changes in Metric Availability

The CPU Demand of Recommended (%) metric is no longer available in vRealize Operations Manager version 6.x. To approximate the metric, create a super metric using the following calculations, and add it to your Views and Reports as needed.

```
( (CPU|Stress Free Demand (MHz)) × (CPU|Current Size in Unit(s)) ) ÷ ( (CPU|Recommended Size (vCPUs)) × (CPU|Current Size (MHz)) )
```

For more information about super metrics, see Configuring Super Metrics .

Metrics for vCenter Server Components

vRealize Operations Manager connects to VMware vCenter Server® instances through the vCenter adapter to collect metrics for vCenter Server components and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

vCenter Server components are listed in the `describe.xml` file for the vCenter adapter. The following example shows sensor metrics for the host system in the `describe.xml` file.

```
<ResourceGroup instanced="false" key="Sensor" nameKey="1350" validation="">
    <ResourceGroup instanced="false" key="fan" nameKey="1351" validation="">
        <ResourceAttribute key="currentValue" nameKey="1360" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal=""  unit="percent"/>
        <ResourceAttribute key="healthState" nameKey="1361" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal="" />
    </ResourceGroup>
    <ResourceGroup instanced="false" key="temperature" nameKey="1352" validation="">
        <ResourceAttribute key="currentValue" nameKey="1362" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal="" />
        <ResourceAttribute key="healthState" nameKey="1363" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal="" />
    </ResourceGroup>
 </ResourceGroup>
```

Each `ResourceAttribute` element includes the name of a metric that appears in the UI and is documented as a Metric Key.

Table 7-1. Sensor Metrics for Host System Cooling

Metric Key	Metric Name	Description
Sensor|fan|currentValue	Speed	Fan speed.
Sensor|fan|healthState	Health State	Fan health state.
Sensor|temperature|currentValue	Temperature	Host system temperature.
Sensor|temperature|healthState	Health State	Host system health state.

vSphere Metrics

vRealize Operations Manager collects CPU use, disk, memory, network, and summary metrics for objects in the vSphere world.

Capacity metrics can be calculated for vSphere world objects. See Capacity and Project-Based Metrics.

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Table 7-2. CPU Usage Metrics

Metric Name	Description
CPU|Capacity usage	CPU usages as a percent during the interval. Key: cpu|capacity_usagepct_average
CPU|CPU contention(%)	This metric shows the percentage of time the VMs in the ESXi hosts are unable to run because they are contending for access to the physical CPUs. The number shown is the average number for all VMs. The number will be lower than the highest number experienced by the VM that is most impacted by CPU contention. Use this metric to verify if the host can serve all its VMs efficiently. Low contention means that the VM can access everything it demands to run smoothly. It means that the infrastructure is providing good service to the application team. When using this metric, ensure that the number is within your expectation. Look at both the relative number and the absolute number. Relative means a drastic change in value, meaning that the ESXi is unable to serve the VMs. Absolute means that the real value itself is high. Investigate why the number is high. One factor that impacts this metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz, the reduction in speed is accounted for because it shows that the VM is not running at full speed. This metric is calculated in the following way: cpu|capacity_contention / (200 * summary|number_running_vcpus) Key: cpu|capacity_contentionPct
CPU|Demand (%)	This metric shows the amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. This metric represents the average active CPU load for the past five minutes. Keep this number below 100% if you set the power management to maximum. This metric is calculated in the following way: (cpu.demandmhz / cpu.capacity_provisioned)*100 Key: cpu|demandPct
CPU|Demand (MHz)	This metric shows the amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. Key: cpu|demandmhz
CPU|Demand	CPU demand in megahertz. Key: cpu|demand_average
CPU|IO wait	IO wait (ms). Key: cpu|iowait
CPU|number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Provisioned Capacity (MHz)	capacity in MHz of the physical CPU cores. Key: cpu|capacity_provisioned

Table 7-2. CPU Usage Metrics (continued)

Metric Name	Description
CPU|Provisioned vCPU(s)	Number of provisioned CPU cores. Key: cpu|corecount_provisioned
CPU|Reserved Capacity (MHz)	Total CPU capacity reserved by virtual machines. Key: cpu|reservedCapacity_average
CPU|Usage (MHz)	CPU usages, as measured in megahertz, during the interval. ■ VM - Amount of actively used virtual CPU. This is the host's view of the CPU usage, not the guest operating system view. ■ Host - Sum of the actively used CPU of all powered on virtual machines on a host. The maximum possible value is the frequency of the two processors multiplied by the number of processors. For example, if you have a host with four 2 GHz CPUs running a virtual machine that is using 4000 MHz, the host is using two CPUs completely: 400 / (4 2000) = 0.50 Key: cpu|usagemhz_average
CPU|Wait	Total CPU time spent in wait state. The wait total includes time spent in the CPU Idle, CPU Swap Wait, and CPU I/O Wait states. Key:cpu|wait
CPU|Workload (%)	Percent of workload Key: cpu|workload

Memory Metrics

Memory metrics provide information about memory use and allocation.

Table 7-3. Memory Metrics

Metric Name	Description
mem|Contention (%)	This metric shows the percentage of time VMs are waiting to access swapped memory. Use this metric to monitor ESXi memory swapping. A high value indicates that the ESXi is running low on memory, and a large amount of memory is being swapped. Key: mem|host_contentionPct
mem|Machine Demand (KB)	Host memory demand in kilobytes. Key: mem|host_demand
mem|Provisioned Memory	Provisioned host memory in kilobytes. Key: mem|host_provisioned
mem|Reserved Capacity (KB)	Total amount of memory reservation used by powered-on virtual machines and vSphere services on the host. Key: mem|reservedCapacity_average
mem|Usable Memory (KB)	Usable host memory in kilobytes. Key: mem|host_usable
mem|Host Usage (KB)	Host memory use in kilobytes. Key: mem|host_usage

Table 7-3. Memory Metrics (continued)

Metric Name	Description
mem|Usage/Usable (%)	Memory usage as percentage of total configured or available memory. Key: mem|host_usagePct
mem|Workload (%)	Percent of workload. Key: mem|workload

Network Metrics

Network metrics provide information about network performance.

Table 7-4. Network Metrics

Metric Name	Description
net|Packets Dropped (%)	This metric shoes the percentage of received and transmitted packets dropped in the collection interval. Use this metric to monitor the reliability and performance of the ESXi network. A high value indicates that the network is not reliable and performance decreases. Key: net|droppedPct
net|Usage Rate (KB per second)	Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine. Key: net|usage_average
net|Workload (%)	Percent of workload. Key: net|workload

Disk Metrics

Disk metrics provide information about disk use.

Table 7-5. Disk Metrics

Metric Name	Description
disk|Commands per second	Average number of commands issued per second during the collection cycle. Key: disk|commandsAveraged_average
disk|Usage Rate (KB per second)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. Key: disk|usage_average
disk|Workload (%)	Percent of workload. Key: disk|workload

Summary Metrics

Summary metrics provide information about overall performance.

Table 7-6. Summary Metrics

Metric Name	Description
summary|Number of Running Hosts	Number of running hosts. Key: summary|number_running_hosts
summary|Number of Running VMs	This metric shows the number of running VMs at a given point in time. The data is sampled every five minutes. A large number of running VMs might be a reason for CPU or memory spikes because more resources are used in the host. The number of running VMs gives you a good indicator of how many requests the ESXi host must juggle. Powered off VMs are not included because they do not impact ESXi performance. A change in the number of running VMs can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, because all the VMs will fail if the ESXi crashes. Use this metric to look for a correlation between spikes in the running VMs and spikes in other metrics such as CPU contention, or memory contention. Key: summary|number_running_vms
summary|Total Number of Clusters	Total number of clusters. Key: summary|total_number_clusters
summary|Total Number of Datastores	Total number of datastores. Key: summary|total_number_datastores
summary|Total Number of Hosts	Total number of hosts. Key: summary|total_number_hosts
summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
summary|Total Number of Datacenters	Total number of data centers. Key: summary|total_number_datacenters
summary|Number VCPUs on Powered on VMs	Number of virtual CPUs on powered-on virtual machines. Key: summary|number_running_vcpus
summary|Average Running VM Count per Running Host	Average running virtual machine count per running host. Key: summary|avg_vm_density

vCenter Server Metrics

vRealize Operations Manager collects CPU use, disk, memory, network, and summary metrics for vCenter Server system objects.

vCenter Server metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Table 7-7. CPU Usage Metrics

Metric Key	Metric Name	Description
cpu|capacity_usagepct_average	Capacity Usage (%)	Percent capacity used.
cpu|capacity_contentionPct	CPU Contention (%)	Percent CPU contention.
cpu|demandPct	Demand (%)	Percent demand.
cpu|demandmhz	Demand (MHz)	Demand in megahertz.
cpu|demand_average	Demand	CPU Demand.
cpu|iowait	IO Wait (ms)	IO wait time in milliseconds.
cpu|numpackages	Number of CPU Sockets	Number of CPU sockets.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|capacity_provisioned	Provisioned Capacity (MHz)	Provisioned capacity in megahertz.
cpu|corecount_provisioned	Provisioned vCPU	Number of provisioned virtual CPU cores.
cpu|reservedCapacity_average	Reserved Capacity (MHz)	Sum of the reservation properties of the immediate children of the host's root resource pool.
cpu|usagemhz_average	Usage (MHz)	Average CPU use in megahertz.
cpu|wait	Wait (ms)	CPU time spent on the idle state.
cpu|overhead_average	Overhead	Amount of CPU that is overhead.
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|vm_capacity_provisioned	Provisioned Capacity	Provisioned capacity (MHz).

Datastore Metrics

Datastore metrics provide information about the datastore.

Table 7-8. Datastore Metrics

Metric Key	Metric Name	Description
datastore|maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore|maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Maximum observed number of outstanding IO operations.
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore|numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.

Table 7-8. Datastore Metrics (continued)

Metric Key	Metric Name	Description
datastore|numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.

Disk Metrics

Disk metrics provide information about disk use.

Table 7-9. Disk Metrics

Metric Key	Metric Name	Description
disk|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection cycle.
disk|totalLatency_average	Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Device Command Latency and Physical Device Command Latency metrics.
disk|usage_average	Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.
disk|sum_queued_oio	Total queued outstanding operations	Sum of queued operations and outstanding operations.
disk|max_observed	Max Observed OIO	Max observed IO for a disk.

Diskspace Metrics

Disk space metrics provide information about disk space use.

Table 7-10. Diskspace Metrics

Metric Key	Metric Name	Description
diskspace|total_usage	Total disk space used (KB)	Total disk space used on all datastores visible to this object.
diskspace|total_capacity	Total disk space (KB)	Total disk space on all datastores visible to this object.
diskspace|total_provisioned	Total provisioned disk space (KB)	Total provisioned disk space on all datastores visible to this object.

Memory Metrics

Memory metrics provide information about memory use and allocation.

Table 7-11. Memory Metrics

Metric Key	Metric Name	Description
mem|host_contentionPct	Contention (%)	Percent host memory contention.
mem|host_demand	Machine Demand (KB)	Host memory demand in kilobytes.
mem|host_systemUsage	ESX System Usage	Memory usage by the VMkernel and ESX user-level services.
mem|host_provisioned	Provisioned Memory (KB)	Provisioned host memory in kilobytes.
mem|reservedCapacity_average	Reserved Capacity (KB)	Sum of the reservation properties of the immediate children of the host's root resource pool.
mem|host_usable	Usable Memory (KB)	Usable host memory in kilobytes.
mem|host_usage	Host Usage (KB)	Host memory use in kilobytes.
mem|host_usagePct	Usage/Usable (%)	Percent host memory used.
mem|host_contention	Contention (KB)	Host contention in kilobytes.
mem|overhead_average	VM Overhead (KB)	Memory overhead reported by host.

Network Metrics

Network metrics provide information about network performance.

Table 7-12. Network Metrics

Metric Key	Metric Name	Description
net|droppedPct	Packets Dropped (%)	Percent network packets dropped.
net|usage_average	Usage Rate (KBps)	Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine.
net|packetsRx_summation	Packets Received	Number of packets received in the performance interval.
net|packetsTx_summation	Packets Transmitted	Number of packets transmitted in the performance interval.
net|droppedRx_summation	Received Packets Dropped	Number of received packets dropped in the performance interval.
net|droppedTx_summation	Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval.
net|maxObserved_KBps	Max Observed Throughput (KBps)	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput (KBps)	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput (KBps)	Max observed received rate of network throughput.
net|transmitted_average	Data Transmit Rate (KBps)	Average amount of data transmitted per second.
net|received_average	Data Receive Rate (KBps)	Average amount of data received per second.

Summary Metrics

Summary metrics provide information about overall performance.

Table 7-13. Summary Metrics

Metric Key	Metric Name	Description
summary|number_running_hosts	Number of Running Hosts	Number of hosts that are on.
summary|number_running_vms	Number of Running VMs	Number of virtual machines that are on.
summary|total_number_clusters	Total Number of Clusters	Total number of clusters.
summary|total_number_datastores	Total Number of Datastores	Total number of datastores.
summary|total_number_hosts	Total Number of Hosts	Total number of hosts.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.
summary|workload_indicator	Workload Indicator (%)	Percent workload indicator.
summary|total_number_datacenters	Total Number of Datacenters	Total number of datacenters.
summary|number_powered_on_cores	Number of Cores on Powered On Hosts	Number of cores on powered-on hosts.
summary|number_running_vcpus	Number VCPUs on Powered on VMs	Number of virtual CPUs on powered-on virtual machines.
summary|avg_vm_density	Average Running VM Count per Running Host	Average running virtual machine count per running host.
summary|vc_query_time	VC Query Time (ms)	vCenter Server query time in milliseconds.
summary|derived_metrics_comp_time	Derived Metrics Computation Time (ms)	Derived metrics computation time in milliseconds.
summary|number_objs	Number of objects	Number of objects.
summary|number_vc_events	Number of VC Events	Number of vCenter Server events.
summary|number_sms_metrics	Number of SMS Metrics	Number of SMS metrics.
summary|collector_mem_usage	Collector Memory Usage (MB)	Collector memory use in megabytes.

Virtual Machine Metrics

vRealize Operations Manager collects configuration, CPU use, memory, datastore, disk, virtual disk, guest file system, network, power, disk space, storage, and summary metrics for virtual machine objects.

Capacity metrics can be calculated for virtual machine objects. See Capacity and Project-Based Metrics.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the virtual machines in your environment.

Configuration Metrics for Virtual Machines

Configuration metrics provide information about virtual machine configuration.

Metric	Description
Config|Thin Provisioned Disk	Thin Provisioned Disk. Key: config|hardware|thin_Enabled
Config|Number of CPUs	Number of CPUs for a Virtual Machine. Key: config|hardware|num_Cpu
Config|Disk Space	Disk space metrics. Key: config|hardware|disk_Space

CPU Usage Metrics for Virtual Machines

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|IO Wait (ms)	CPU time spent waiting for IO. Key: cpu|iowait
CPU|Wait (ms)	Wait time in milliseconds. Key: cpu|wait
CPU|Overall CPU Contention (ms)	The amount of time the CPU cannot run due to contention. Key: cpu|capacity_contention
CPU|Reservation Used	CPU Reservation Used. Key: cpu|reservation_used
CPU|Effective Limit	CPU Effective Limit. Key: cpu|effective_limit
CPU|Estimated Entitlement	CPU Estimated Entitlement. key: cpu|estimated_entitlement
CPU|Idle (%)	Percentage time that CPU is idle. Key: cpu|idlePct
CPU|IO Wait (%)	Percentage IO Wait. Key: cpu|iowaitPct
CPU|Swap wait (%)	Percentage swap wait for CPU. Key: cpu|swapwaitPct
CPU|Wait (%)	Percentage of total CPU time spent in wait state. Key: cpu|waitPct
CPU|System (%)	Percentage CPU time spent on system processes. Key: cpu|systemSummationPct
CPU|Demand Over Limit (MHz)	Amount of CPU Demand that is over the configured CPU Limit. Key: cpu|demandOverLimit

Metric	Description
CPU|Demand Over Capacity (MHz)	Amount of CPU Demand that is over the configured CPU Capacity. Key: cpu|demandOverCapacity
CPU|Recommended Size Reduction (%)	Percentage of recommended CPU size reduction. Key: cpu|sizePctReduction
CPU|Normalized Co-stop	Percentage of co-stop time, normalized across all vCPUs. Key: cpu|perCpuCoStopPct
CPU|Recommended number of vCPUs to Add	Recommended number of vCPUs to Add to the VM. Key: cpu|numberToAdd
CPU|Recommended number of vCPUs to Remove	Recommended number of vCPUs to Remove from the VM. Key: cpu|numberToRemove
CPU|Capacity entitlement (MHz)	CPU entitlement for the VM after taking limits into account. Key: cpu|capacity_entitlement
CPU|Provisioned CPU Cores	Number of provisioned CPU cores. Key: cpu|corecount_provisioned
CPU|Capacity Demand Entitlement (%	Percent capacity demand entitlement. Key: cpu|capacity_demandEntitlementPct
* CPU|CPU Contention (%)	CPU contention as a percentage of 20-second collection interval. Key: cpu|capacity_contentionPct
CPU|Capacity Provisioned	Provisioned CPU capacity in megahertrz. Key: cpu|capacity_provisioned
CPU|Demand (MHz)	CPU demand in megahertz. Key: cpu|demandmhz
CPU|Host demand for aggregation	Host demand for aggregation. Key: cpu|host_demand_for_aggregation
CPU|Demand (ms)	The total CPU time that the VM could use if there was no contention. Key: cpu|demand_average
CPU|Demand (%)	CPU demand as a percentage of the provisioned capacity. Key: cpu|demandPct
CPU|Dynamic Entitlement	CPU Dynamic Entitlement. Key: cpu|dynamic_entitlement
* CPU|Usage (%)	This metric indicates the percentage of CPU that was used out of all the CPU that was allocated to the VM. CPU usage can indicate if the VM is undersized. Key: cpu|usage_average

Metric	Description
CPU|Usage (MHz)	CPU use in megahertz. Key: cpu|usagemhz_average
CPU|System (ms)	CPU time spent on system processes. Key: cpu|system_summation
CPU|Wait (ms)	Total time that a virtual CPU can not be run. It could be idle (halted) or waiting for an external event such as I/O. Key: cpu|wait_summation
CPU|Ready (ms)	CPU time spent in the ready state. Key: cpu|ready_summation
* CPU|Ready (%)	This metric indicates the percentage of time in which the VM was waiting in line to use the CPU on the host. A large ready time for a VM indicates that the VM needed CPU resources but the infrastructure was busy serving other VMs. This might indicate that the host is trying to serve too many VMs. Whenever the CPU ready is larger than 10%, you should check if the host is overloaded, or if the VM really needs all the resources that were allocated to it. Key: cpu|readyPct
CPU|Used (ms)	CPU time that is used. Key: cpu|used_summation
CPU|Extra (ms)	Extra CPU time in milliseconds. Key: cpu|extra_summation
CPU|Guaranteed (ms)	CPU time that is guaranteed for the virtual machine. Key: cpu|guaranteed_latest
CPU|Swap Wait (ms)	Swap wait time in milliseconds. Key: cpu|swapwait_summation
CPU|Co-stop (ms)	Time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costop_summation
CPU|Co-stop (%)	Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costopPct
CPU|Idle (ms)	CPU time that is idle. Key: cpu|idle_summation
CPU|Latency	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpu|latency_average
CPU|Max Limited	Time the VM is ready to run, but is not run due to maxing out its CPU limit setting. Key: cpu|maxlimited_summation

Metric	Description
CPU|Overlap	Time the VM was interrupted to perform system services on behalf of that VM or other VMs. Key: cpu|overlap_summation
CPU|Run	Time the VM is scheduled to run. Key: cpu|run_summation
CPU|Entitlement Latest	Entitlement Latest. Key: cpu|entitlement_latest

CPU Utilization for Resources Metrics for Virtual Machines

CPU utilization for resources metrics provide information about resource CPU use.

Metric	Description
rescpu|CPU Active (%) (*interval*)	The average active time (actav) or peak active time (actpk) for the CPU during various intervals. Key: rescpu|actav1_latest rescpu|actav5_latest rescpu|actav15_latest rescpu|actpk1_latest rescpu|actpk5_latest rescpu|actpk15_latest
rescpu|CPU Running (%) (*interval*)	The average runtime (runav) or peak active time (runpk) for the CPU during various intervals. Key: rescpu|runav1_latest rescpu|runav5_latest rescpu|runav15_latest rescpu|runpk1_latest rescpu|runpk5_latest rescpu|runpk15_latest
rescpu|CPU Throttled (%) (*interval*)	Amount of CPU resources over the limit that were refused, average over various intervals. Key: rescpu|maxLimited1_latest rescpu|maxLimited5_latest rescpu|maxLimited15_latest
rescpu|Group CPU Sample Count	The sample CPU count. Key: rescpu|sampleCount_latest
rscpu|Group CPU Sample Period (ms)	The sample period. Key: rscpu|samplePeriod_latest

Memory Metrics for Virtual Machines

Memory metrics provide information about memory use and allocation.

Metric	Description
Mem|Host Active (KB)	Host active memory use in kilobytes. Key: mem|host_active
Mem|Usage (KB)	Memory use in kilobytes. Key: mem|host_usage
Mem|Contention (KB)	Memory contention in kilobytes. Key: mem|host_contention
Mem|Contention (%)	Percent memory contention. Key: mem|host_contentionPct
Mem|Guest Configured Memory (KB)	Guest operating system configured memory in kilobytes. Key: mem|guest_provisioned
Mem|Guest Dynamic Entitlement (KB)	Guest Memory Dynamic Entitlement. Key: mem|guest_dynamic_entitlement
Mem|Guest Active Memory (%)	Percent guest operating system active memory. Key: mem|guest_activePct
Mem|Guest Non Pageable Memory (KB)	Guest operating system non-pageable memory in kilobytes. Key: mem|guest_nonpageable_estimate
Mem|Reservation Used	Memory Reservation Used. Key: mem|reservation_used
Mem|Effective Limit	Memory Effective Limit. Key: mem|effective_limit
Mem|Estimated Entitlement	Memory Estimated Entitlement. Key: mem|estimated_entitlement
Mem|Demand for aggregation	Host demand for aggregation. Key: mem|host_demand_for_aggregation
Mem|NUMA Remote Latest	Non-uniform memory access Remote (Kb). Key: mem|numa.remote_latest
Mem|NUMA Local Latest	Non-uniform memory access Local (Kb). Key: mem|numa.local_latest
Mem|NUMA Migrations Latest	Non-uniform memory access Migrations (number). Key: mem|numa.migrations_latest
Mem|NUMA Locality Average	Non-uniform memory access Locality (%). Key: mem|numa.locality_average
Mem|Demand Over Limit	Amount of Memory Demand that is over the configured Memory Limit. Key: mem|demandOverLimit
Mem|Demand Over Capacity	Amount of Memory Demand that is over the configured Memory Capacity. Key: mem|demandOverCapacity

Metric	Description
Mem|Recommended Size Reduction (%)	Percentage of recommended Memory size reduction. Key: mem|sizePctReduction
Mem|Balloon (%)	Percentage of total memory that has been reclaimed via ballooning. Key: mem|balloonPct
* Mem|Guest Usage (KB)	This metric shows the amount of memory the VM uses. Key: mem|guest_usage
Mem|Guest Demand (KB)	Guest operating system demand in kilobytes. Key: mem|guest_demand
Mem|Guest Non Pageable Memory (KB)	Guest operating system non-pageable memory in kilobytes. Key: mem|host_nonpageable_estimate
Mem|Host Demand (KB)	Memory demand in kilobytes. Key mem|host_demand
Mem|Demand with Reservation (KB)	Memory Demand with Reservation considered in KB. Key: mem|host_demand_reservation
Mem|Guest Workload	Guest Workload (%). Key: mem|guest_workload
Mem|Host Workload	Host Workload (%). Key: host_workload
Mem|Balloon (%)	Amount of memory currently used by the virtual machine memory control. Key: mem|vmmemctl_average
Mem|Guest Active (%)	Amount of memory that is actively used. Key: mem|active_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Shared (KB)	Amount of shared memory in kilobytes. Key: mem|shared_average
Mem|Zero (KB)	Amount of memory that is all 0. Key: mem|zero_average
* Mem|Swapped (KB)	This metric shows how much memory is being swapped. Meaning, the amount of unreserved memory in kilobytes. Key: mem|swapped_average
Mem|Swap Target (KB)	Amount of memory that can be swapped in kilobytes. Key: mem|swaptarget_average
Mem|Swap In (KB)	Swap-in memory in kilobytes. Key: mem|swapin_average
Mem|Swap Out (KB)	amount of memory swapped out in kilobytes. Key: mem|swapout_average

Metric	Description
* Mem|Usage (%)	This metric shows how much memory is being used out of the allocated memory for the VM. Key: mem|usage_average
Mem|Balloon Target (KB)	Amount of memory that can be used by the virtual machine memory control. Key: mem|vmmemctltarget_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory in kilobytes. Key: mem|consumed_average
Mem|Overhead (KB)	Memory overhead in kilobytes. Key: mem|overhead_average
Mem|Host Dynamic Entitlement	Mem Machine Dynamic Entitlement. Key: mem|host_dynamic_entitlement
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem|swapinRate_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory to disk during the current interval. Key: mem|swapoutRate_average
Mem|Active Write (KB)	Active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Compressed memory in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Compression rate in kilobytes per second. Key: mem|commpressionRate_average
Mem|Decompression Rate (KBps)	Decompression rate in kilobytes per second. Key: mem|decompressionRate_average
Mem|Overhead Max (KB)	Maximum overhead in kilobytes. Key: mem|overheadMax_average
Mem|Zip Saved (KB)	Zip-saved memory in kilobytes. Key: mem|zipSaved_latest
Mem|Zipped (KB)	Zipped memory in kilobytes. Key: mem|zipped_latest
Mem|Entitlement	Amount of host physical memory the VM is entitled to, as determined by the ESX schedule. Key: mem|entitlement_average
Mem|Latency	Percentage of time the VM is waiting to access swapped or compressed memory. Key: mem|latency_average
Mem|Capacity Contention	Capacity Contention. Key: mem|capacity.contention_average

Metric	Description
Mem|Swap In Rate from Host Cache	Rate at which memory is being swapped from host cache into active memory. Key: mem|llSwapInRate_average
Mem|Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem|llSwapOutRate_average
Mem|Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem|llSwapUsed_average
Mem|Overhead Touched	Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM. Key: mem|overheadTouched_average

Datastore Metrics for Virtual Machines

Datastore metrics provide information about datastore use.

Metric	Description
Datastore|Commands per second	Average number of commands issued per second during the collection interval. Key:datastore|commandsAveraged_average
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore|oio
Datastore|Demand	Datastore demand. Key: datastore|demand
Datastore|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Usage Average (KBps). Key: datastore|usage_average
Datastore|Used Space (MB)	Used space in megabytes. Key: datastore|used
Datastore|Not Shared (GB)	Space used by VMs that is not shared. Key: datastore|notshared
* Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
* Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average

Metric	Description
* Datastore|Read Rate (KBps)	This metric shows the amount of data that the VM reads to the datastore per second.. Key: datastore|read_average
*Datastore|Read Latency (ms)	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
*Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average
* Datastore|Write Rate	This metric shows the amount of data that the VM writes to the datastore per second. Key: datastore|write_average
Datastore|Highest Latency	Highest Latency. Key: datastore|maxTotalLatency_latest
Datastore|Total Latency Max	Total Latency Max (ms). Key: datastore|totalLatency_max
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO

Disk Metrics for Virtual Machines

Disk metrics provide information about disk use.

Metric	Description
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_average
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average

Metric	Description
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Usage Rate (KBps)	Use rate in kilobytes per second. Key: disk|usage_average
Disk|I/O Usage Capacity	This metric is a function of storage|usage_average and disk|workload. storage|usage_average is an average over all storage devices. This means that disk|usage_capacity is not specific to the selected VM or the host of the VM. Key: disk|usage_capacity
Disk|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: disk|diskoio
Disk|Queued Operations	Queued operations. Key: disk|diskqueued
Disk|Demand (%)	Percent demand. Key: disk|diskdemand
Disk |Total Queued Outstanding Operations	Sum of Queued Operation and Outstanding Operations. Key: disk |sum_queued_oio
Disk|Max Observed OIO	Max Observed IO for a disk. Key: disk|max_observed
Disk|Read Rate (KBps)	Amount of data read in the performance interval. Key: disk|read_average
Disk|Write Rate (KBps)	Amount of data written to disk in the performance interval. Key: disk|write_average
Disk|Read Requests	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Write Requests	Number of times data was written to the disk in the defined interval. Key: disk|numberWrite_summation
Disk|Bus Resets	The number of bus resets in the performance interval. Key: disk|busResets_summation
Disk|Commands Issued	The number of disk commands issued in the performance interval. Key: disk|commands_summation
Disk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: disk|commandsAborted_summation
Disk|Highest Latency	Highest latency. Key: disk|maxTotalLatency_latest

Metric	Description
Disk|SCSI Reservation Conflicts	SCSI Reservation Conflicts. Key: disk|scsiReservationConflicts_summation
Disk|Disk Read Latency	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalReadLatency_average
Disk|Disk Write Latency	The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency. Key: disk|totalWriteLatency_average
Disk|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: disk|totalLatency_average

Virtual Disk Metrics for Virtual Machines

Virtual disk metrics provide information about virtual disk use.

Metric	Description
VirtualDisk|Usage	Average CPU usage as a percentage. Key: virtualDisk|usage
VirtualDisk|Total Latency	Total latency. Key: virtualDisk|totalLatency
VirtualDisk|Commands Per Second	Average number of commands per second. Key: virtualDisk|commandsAveraged_average
VirtualDisk|Read Requests	Average number of read commands issued per second to the virtual disk during the collection interval. Key: virtualDisk|numberReadAveraged_average
VirtualDisk|Write Requests	Average number of write commands issued per second to the virtual disk during the collection interval. Key: virtualDisk|numberWriteAveraged_average
VirtualDisk|Read Rate (KBps)	Rate of reading data from the virtual disk in kilobytes per second. Key: virtualDisk|read_average
VirtualDisk|Read Latency (ms)	Average amount of time for a read operation from the virtual disk. Total latency = kernel latency + device latency. Key: virtualDisk|totalReadLatency_average
VirtualDisk|Write Latency (ms)	Average amount of time for a write operation to the virtual disk. Total latency = kernel latency + device latency. Key: virtualDisk|totalWriteLatency_average

Metric	Description
VirtualDisk|Write Rate (KBps)	Rate of writing data from the virtual disk in kilobytes per second. Key: virtualDisk|write_average
VirtualDisk|Bus Resets	The number of bus resets in the performance interval. Key: virtualDisk|busResets_summation
VirtualDisk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: virtualDisk|commandsAborted_summation
VirtualDisk|Read Load	Storage DRS virtual disk metric read load. Key: virtualDisk|readLoadMetric_latest
VirtualDisk|Outstanding Read Requests	Average number of outstanding read requests to the virtual disk. Key: virtualDisk|readOIO_latest
VirtualDisk|Write Load	Storage DRS virtual disk write load. Key: virtualDisk|writeLoadMetric_latest
VirtualDisk|Outstanding Write Requests	Average number of outstanding write requests to the virtual disk. Key: virtualDisk|writeOIO_latest
VirtualDisk|Number of Small Seeks	Small Seeks. Key: virtualDisk|smallSeeks_latest
VirtualDisk|Number of Medium Seeks	Medium Seeks. Key: virtualDisk|mediumSeeks_latest
VirtualDisk|Number of Large Seeks	Large Seeks. Key: virtualDisk|largeSeeks_latest
VirtualDisk|Read Latency (microseconds)	Read Latency in microseconds. Key: virtualDisk|readLatencyUS_latest
VirtualDisk|Write Latency (microseconds)	Write Latency in microseconds. Key: virtualDisk|writeLatencyUS_latest
VirtualDisk|Average Read request size	Read IO size. Key: virtualDisk|readIOSize_latest
VirtualDisk|Average Write request size	Write IO size. Key: virtualDisk|writeIOSize_latest

Guest File System Metrics for Virtual Machines

Guest file system metrics provide information about guest file system capacity and free space.

Metric	Description
Guest file system|Guest File System Capacity (MB)	Total capacity on guest file system in megabytes. Key: guestfilesystem|capacity
Guest file system|Guest File System Free (MB)	Total free space on guest file system in megabytes. Key: guestfilesystem|freespace

Metric	Description
Guest file system|Guest File System Usage (%)	Percent guest file system. Key: guestfilesystem|percentage
Guest file system|Guest File System Usage	Total usage of guest file system. Key: guestfilesystem|usage
* Guest file system|Total Guest File System Free (GB)	This metric displays the amount of free disk space from all file systems attached to this VM. Use this metric to see if there are spikes in the free space or if there is an organic growth for this VM. Key: guestfilesystem|freespace_total
* Guest file system|Total Guest File System Capacity (GB)	This metric displays the amount of disk space allocated for the VM. Correlate other metrics with this metric to indicate if changes occur in the disk space allocation for the VM. Key: guestfilesystem|capacity_total
* Guest file system|Total Guest File System Usage (%)	This metric displays the amount of displa space being used out of the total allocated disk space. Use his metric to track if the overall usage is stable, or if it reaches the limits. You should avoid having VMs with a disk space usage of >95% since this could impact your system. Key: guestfilesystem|percentage_total
Guest file system|Total Guest File System Usage	Total usage of guest file system. Key: guestfilesystem|usage_total

Network Metrics for Virtual Machines

Network metrics provide information about network performance.

Metric	Description
Net|Demand (%)	Percent demand. Key: net|demand
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|Packets Received per second	Number of packets received in the performance interval. Key: net|packetsRxPerSec
Net|Packets Transmitted per second	Number of packets transmitted in the performance interval. Key: net|packetsTxPerSec
* Net|Data Transmit Rate (KBps)	This metric shows the rate of data being sent by the VM per second. Key: net|transmitted_average
* Net|Data Receive Rate (KBps)	This metric shows the rate of data received by the VM per second. Key: net|received_average

Metric	Description
Net|Packets per second	Number of packets transmitted and received per second. Key: net|PacketsPerSec
Net|I/O Usage Capacity	IO use capacity. Key: net|usage_capacity
Net|Max Observed Throughput (KBps)	Maximum observed throughput in kilobytes per second. Key: net|maxObserved_KBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
* Net|Received Packets Dropped	This metric shows the number of received packets dropped in the collection interval. Key: net|droppedRx_summation
* Net|Transmitted Packets Dropped	This metric shows the number of transmitted packets dropped in the collection interval Key: net|dropppedTx_summation
Net|Packets Dropped (%)	Percentage of packets dropped. Key: net|droppedPct
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|Broadcast Packets Transmitted	Number of broadcast packets transmitted during the sampling interval. Key: net|broadcastTx_summation
Net|Broadcast Packets Received	Number of broadcast packets received during the sampling interval. Key: net|broadcastRx_summation
Net|Bytes Rx (KBps)	Average amount of data received per second. Key: net|bytesRx_average
Net|Bytes Tx (KBps)	Average amount of data transmitted per second. Key: net|bytesTx_average
Net|Multicast Packets Received	Number of multicast packets received. Key: net|multicastRx_summation
Net|Multicast Packets Transmitted	Number of multicast packets transmitted. Key: net|multicastTx_summation
Net|VM to Host Data Transmit Rate	Average amount of data transmitted per second between VM and host. Key: net|host_transmitted_average

Metric	Description
Net|VM to Host Data Receive Rate	Average amount of data received per second between VM and host. Key: net|host_received_average
Net|VM to Host Usage Rate	The sum of the data transmitted and received for all the NIC instances between VM and host. Key: net|host_usage_average
Net|VM to Host Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput between VM and host. Key: net|host_maxObserved_Tx_KBps
Net|VM to Host Max Observed Received Throughput	Max observed received rate of network throughput between VM and host. Key: net|host_maxObserved_Rx_KBps
Net|VM to Host Max Observed Throughput	Max observed rate of network throughput between VM and host. Key: net|host_maxObserved_KBps
Net|Data Transmit Demand Rate	Data transmitted demand rate. Key: net|transmit_demand_average
Net|Data Receive Demand Rate	Data received demand rate. Key: net|receive_demand_average

System Metrics for Virtual Machines

System metrics for virtual machines provide general information about the virtual machine, such as its build number and running state.

Metric	Description
Sys|Powered ON	Powered on virtual machines. 1 if powered on, 0 if powered off, -1 if unknown Key: sys|poweredOn
Sys|Uptime (seconds)	Number of seconds since system startup. Key: sys|uptime_latest
Sys|Heartbeat	Number of heartbeats from the virtual machine in the defined interval. Key: sys|heartbeat_summation
Sys|vMotion Enabled	1 if vMotion is enabled or 0 if vMotion is not enabled. Key: sys|vmotionEnabled
Sys|Product String	VMware product string. Key: sys|productString
Sys|Build Number	VMware build number. Key: sys|build
Sys|OS Uptime	Total time elapsed, in seconds, since last operating system boot-up. Key: sys|osUptime_latest

Power Metrics for Virtual Machines

Power metrics provide information about power use.

Metric	Description
Power|Energy (Joule)	Energy use in joules. Key: power|energy_summation
Power|Power (Watt)	Average power use in watts. Key: power|power_average

Disk Space Metrics for Virtual Machines

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared space in kilobytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disks	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Provisioned Space (GB)	Provisioned space in gigabytes. Key: diskspace|provisioned
Diskspace|Provisioned Space for VM	Provisioned space for VM. Key: diskspace|provisionedSpace
Diskspace|Shared Used (GB)	Shared used space in gigabytes. Key: diskspace|shared
Diskspace|Snapshot Space (GB)	Space used by snapshots. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Space used by virtual disks in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Space used by virtual machine files in gigabytes. Key: diskspace|used
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object.Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned
Diskspace|Active not shared	Unshared disk space used by VMs excluding snapshot. Key: diskspace|activeNotShared

Storage Metrics for Virtual Machines

Storage metrics provide information about storage use.

Metric	Description
Storage|Commands per second	Average number of commands issued per second during the collection interval. Key: storage|commandsAveraged_average
Storage|Contention (%)	Percent contention. Key: storage|contention
Storage|Demand (KBps)	Demand in kilobytes per second. Key: storage|demandKBps
* Storage|Read Latency (ms)	This metric shows the latency that the VM experiences while performing a read action. Key: storage|totalReadLatency_average
Storage|Read Rate (KBps)	Read throughput rate in kilobytes per second. Key: storage|read_average
Storage|Reads per second	Average number of read commands issued per second during the collection interval. Key: storage|numberReadAveraged_average
Storage|Total Latency (ms)	Total latency in milliseconds. Key: storage|totalLatency_average
Storage|Total Usage (KBps)	Total throughput rate in kilobytes per second. Key: storage|usage_average
* Storage|Write Latency (ms)	This metric shows the latency that this VM experiences while performing a write action. Key: storage|totalWriteLatency_average
Storage|Write Rate (KBps)	Write throughput rate in kilobytes per second. Key: storage|write_average
Storage|Writes per second	Average number of write commands issued per second during the collection interval. Key: storage|numberWriteAveraged_average

Summary Metrics for Virtual Machines

Summary metrics provide information about overall performance.

Metric	Description
Summary|Workload Indicator (%)	Percent workload indicator. Key: summary|workload_indicator
Summary|CPU Shares	CPU shares. Key: summary|cpu_shares
Summary|Memory Shares	Memory shares. Key: summary|mem_shares
Summary|Number of Datastores	Number of datastores. Key: summary|number_datastore

Metric	Description
Summary|Number of Networks	Number of networks. Key: summary|number_network
Summary|Running	Number of running virtual machines. Key: summary|running
Summary|Desktop Status	Horizon View Desktop Status. Key: summary|desktop_status

Host System Metrics

vRealize Operations Manager collects many metrics for host systems, including CPU use, datastore, disk, memory, network, storage, and summary metrics for host system objects.

Capacity metrics can be calculated for host system objects. See Capacity and Project-Based Metrics.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the hosts in your environment.

vFlash Module Metrics for Host Systems

vFlash Module metrics provide information about the host system's flash devices.

Metric	Description
vFlashModule|Latest Number of Active VM Disks	Latest Number of Active VM Disks. Key: vflashModule|numActiveVMDKs_latest

Configuration Metrics for Host Systems

Configuration metrics provide information about host system configuration.

Metric	Description
Configuration|Failover Hosts	Failover Hosts. Key: configuration|dasConfig|admissionControlPolicy|failoverHost

Hardware Metrics for Host Systems

Hardware metrics provide information about host system hardware.

Metric	Description
Hardware|Number of CPUs	Number of CPUs for a host. Key: hardware|cpuinfo|num_CpuCores

CPU Usage Metrics for Host Systems

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|Capacity Usage (%)	Percent CPU capacity used. Key: cpu|capacity_usagepct_average
CPU|Usage (%)	Average CPU usage as a percentage. Key: cpu|usage_average
* CPU|CPU Contention (%)	This metric indicates the percentage of time the virtual machines in the ESXi hosts are unable to run because they are contending for access to the physical CPU(s). This is the average number of all VMs. Naturally, the number will be lower than the highest number experienced by the worst hit VM (a VM that suffers the highest CPU contention). Use this metric to verify if the host is able to serve all of its VMs well. When using this metric, ensure the number is within your expectation. The metric is affected by several factors so you need to watch both relative numbers and absolute numbers. Relative means a drastic change in value. This indicates that the ESXi is unable to service its VMs. Absolute means that the real value is high. You should investigate why the value is high. One factor that impacts the CPU contention metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz that reduction in speed is taken into account. This is because the VM is not running at full speed. Key: cpu|capacity_contentionPct
* CPU|Demand (%)	This metric shows the percentage of CPU resources all the VMs would use if there was no CPU contention or any CPU limits set. It represents the average active CPU load for the past five minutes. Keep the number of this metric below 100% if you set Power Management to Maximum. Key: cpu|demandPct
CPU|Demand (MHz)	CPU demand in megahertz. Key: cpu|demandmhz
CPU|IO Wait (ms)	IO wait time in milliseconds. Key: cpu|iowait
CPU|Number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention (ms)	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Provisioned Capacity (MHz)	Capacity in MHz of the physical CPU cores. Key: cpu|capacity_provisioned
CPU|Provisioned virtual CPUs	Provisioned virtual CPUs. Key: cpu|corecount_provisioned

Metric	Description
CPU|Total Wait	CPU time spent in idle state. Key: cpu|wait
CPU|Demand	CPU demand. Key: cpu|demand_average
CPU|Used (msec)	Time accounted to the virtual machine. If a system service runs on behalf of this virtual machine, the time spent by that service (represented by cpu.system) should be charged to this virtual machine. If not, the time spent (represented by cpu.overlap) should not be charged against this virtual machine. Key: cpu|used_summation
CPU|Usage (MHz)	CPU use in megahertz. Key: cpu|usagemhz_average
CPU|Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool. Key: cpu|reservedCapacity_average
CPU|Total Capacity (MHz)	Total CPU capacity in megahertz. Key: cpu|totalCapacity_average
CPU|Idle (ms)	CPU idle time in milliseconds. Key: cpu|idle_summation
CPU|Overhead (KB)	Amount of CPU overhead. Key: cpu|overhead_average
CPU|Demand without overhead	Value of demand excluding any overhead. Key: cpu|demand_without_overhead
CPU|Core Utilization (%)	Percent core utilization. Key: cpu|coreUtilization_average
CPU|Utilization(%)	Percent CPU utilization. Key: cpu|utilization_average
CPU|Core Utilization (%)	Core Utilization. Key: cpu|coreUtilization_average
CPU|Utilization (%)	Utilization. Key: cpu|utilization_average
CPU|Co-stop (ms)	Time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costop_summation
CPU|Latency (%)	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpu|latency_average
CPU|Ready (ms)	Time spent in ready state. Key: cpu|ready_summation
CPU|Run (ms)	Time the virtual machine is scheduled to run. Key: cpu|run_summation

Metric	Description
CPU|Swap wait (ms)	Amount of time waiting for swap space. Key: cpu|swapwait_summation
CPU|Wait (ms)	Total CPU time spent in wait state. Key: cpu|wait_summation
CPU|Provisioned Capacity	Provisioned capacity (MHz). Key: cpu|vm_capacity_provisioned
CPU|Active Host Load For Balance (Long Term)	Active Host Load For Balance (Long Term). Key: cpu|acvmWorkloadDisparityPcttive_longterm_load
CPU|Active Host Load For Balance (Short Term)	Active Host Load For Balance (Short Term). Key: cpu|active_shortterm_load

CPU Utilization for Resources Metrics for Host Systems

CPU utilization for resources metrics provide information about CPU activity.

Metric Name	Description
Rescpu|CPU Active (%) (*interval*)	Average active time for the CPU over the past minute, past five minutes, and at one-minute, five-minute, and 15-minute peak active times. Key: rescpu|actav1_latest rescpu|actav5_latest rescpu|actav15_latest rescpu|actpk1_latest rescpu|actpk5_latest rescpu|actpk15_latest
Rescpu|CPU Running (%) (*interval*)	Average run time for the CPU over the past minute, past five minutes, past 15 minutes, and at one-minute, five-minute, and 15-minute peak times. Key: rescpu|runav1_latest rescpu|runav5_latest rescpu|runav15_latest rescpu|runpk1_latest rescpu|runpk5_latest rescpu|runpk15_latest
Rescpu|CPU Throttled (%) (*interval*)	Scheduling limit over the past minute, past five minutes, and past 15 minutes. Key: rescpu|maxLimited1_latest rescpu|maxLimited5_latest rescpu|maxLimited15_latest

Metric Name	Description
Rescpu|Group CPU Sample Count	Group CPU sample count. Key: rescpu|sampleCount_latest
Rescpu|Group CPU Sample Period (ms)	Group CPU sample period in milliseconds. Key: rescpu|samplePeriod_latest

Datastore Metrics for Host Systems

Datastore metrics provide information about datastore use.

Metric	Notes
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO
Datastore|Commands Averaged	Average number of commands issued per second during the collection interval. Key: datastore|commandsAveraged_average
Datastore|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore|oio
Datastore|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Usage Average (KBps). Key: datastore|usage_average
Datastore|Demand	Demand. Key: datastore|demand
Datastore|Storage I/O Control aggregated IOPS	Aggregate number of IO operations on the datastore. Key: datastore|datastoreIops_average

Metric	Notes
Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average
Datastore|Read Rate (KBps)	Rate of reading data from the datastore in kilobytes per second. Key: datastore|read_average
Datastore|Storage I/O Control normalized latency (ms)	Normalized latency in microseconds on the datastore. Data for all virtual machines is combined. Key: datastore| sizeNormalizedDatastoreLatency_average
Datastore|Read Latency (ms)	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average
Datastore|Write Rate (KBps)	Rate of writing data to the datastore in kilobytes per second. Key: datastore|write_average
Datastore|Max Queue Depth	Max Queue Depth. Key: datastore|datastoreMaxQueueDepth_latest
Datastore|Highest Latency	Highest Latency. Key: datastore|maxTotalLatency_latest
Datastore|Total Latency Max	Total Latency Max (ms). Key: datastore|totalLatency_max
Datastore|Read Latency	Read Latency. Key: datastore|datastoreNormalReadLatency_latest
Datastore|Write Latency	Write Latency. Key: datastore|datastoreNormalWriteLatency_latest
Datastore|Data Read	Data Read. Key: datastore|datastoreReadBytes_latest
Datastore|Data Read Rate	Data Rate. Key: datastore|datastoreReadIops_latest
Datastore|Read Load	Storage DRS metric read load. Key: datastore|datastoreReadLoadMetric_latest
Datastore|Outstanding Read Requests	Outstanding Read Requests. Key: datastore|datastoreReadOIO_latest

Metric	Notes
Datastore|Data Written	Data Written. Key: datastore|datastoreWriteBytes_latest
Datastore|Data Write Rate	Data Write Rate. Key: datastore|datastoreWritelops_latest
Datastore|Write Load	Storage DRS metric write load. Key: datastore|datastoreWriteLoadMetric_latest
Datastore|Outstanding Write Requests	Outstanding Write Requests. Key: datastore|datastoreWriteOIO_latest
* Datastore|Average Observed Virtual Machine Disk I/O Workload	Average Observed Virtual Machine Disk I/O Workload on the Host. Key: datastore|vmPopulationAvgWorkload
Datastore|Maximum Observed VM Disk I/O Workload	Maximum Observed VM Disk I/O Workload on the Host. Key: datastore|vmPopulationMaxWorkload
Datastore|VM Disk I/O Workload Disparity	Percentage Disk I/O workload disparity among the VMs on the Host. Key: datastore|vmWorkloadDisparityPc

Disk Metrics for Host Systems

Disk metrics provide information about disk use.

Metric	Description
Disk|Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. disk|usage_average
Disk|I/O Usage Capacity	This metric is a function of storage|usage_average and disk|workload. storage|usage_average is an average over all storage devices. This means that disk|usage_capacity is not specific to the selected VM or the host of the VM. Key: disk|usage_capacity
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: disk|totalLatency_average
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_average

Metric	Description
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average
Disk|Read Requests	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Write Requests	Number of times data was written to the disk in the defined interval. Key: disk|numberWrite_summation
Disk|Read Rate	Amount of data read in the performance interval. Key: disk|read_average
Disk|Write Rate	Amount of data written to disk in the performance interval. Key: disk|write_average
Disk|Bus Resets	The number of bus resets in the performance interval. Key: disk|busResets_summation
Disk|Commands Issued	The number of disk commands issued in the performance interval. Key: disk|commands_summation
Disk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: disk|commandsAborted_summation
Disk|Physical Device Read Latency (ms)	The average time taken to complete a read from the physical device. Key: disk|deviceReadLatency_average
Disk|Kernel Disk Read Latency (ms)	The average time spent in ESX Server VMKernel per read. Key: disk|kernelReadLatency_average
Disk|Disk Read Latency (ms)	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalReadLatency_average
Disk|Queue Read Latency (ms)	The average time spent in the ESX Server VMKernel queue per read. Key: disk|queueReadLatency_average
Disk|Physical Device Write Latency (ms)	The average time taken to complete a write from the physical device. Key: disk|deviceWriteLatency_average
Disk|Kernel Disk Write Latency (ms)	The average time spent in ESX Server VMKernel per write. Key: disk|kernelWriteLatency_average

Metric	Description
Disk|Disk Write Latency (ms)	The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency. Key: disk|totalWriteLatency_average
Disk|Queue Write Latency (ms)	The average time spent in the ESX Server VMKernel queue per write. Key: disk|queueWriteLatency_average
Disk|Physical Device Command Latency (ms)	The average time taken to complete a command from the physical device. Key: disk|deviceLatency_average
Disk|Kernel Disk Command Latency (ms)	The average time spent in ESX Server VMKernel per command. Key: disk|kernelLatency_average
Disk|Queue Command Latency (ms)	The average time spent in the ESX Server VMKernel queue per command. Key: disk|queueLatency_average
Disk|Number of Outstanding IO Operations	Number of Outstanding IO Operations. Key: disk|diskoio
Disk|Queued Operations	Queued Operations. Key: disk|diskqueued
Disk|Demand	Demand. Key: disk|diskdemand
Disk|Total Queued Outstanding operations	Sum of Queued Operation and Outstanding Operations. Key: disk|sum_queued_oio
Disk|Max Observed OIO	Max Observed IO for a disk. Key: disk|max_observed
Disk|Highest Latency	Highest Latency. Key: disk|maxTotalLatency_latest
Disk|Max Queue Depth	Maximum queue depth during the collection interval. Key: disk|maxQueueDepth_average
Disk|SCSI Reservation Conflicts	SCSI Reservation Conflicts. Key: disk|scsiReservationConflicts_summation

Memory Metrics for Host Systems

Memory metrics provide information about memory use and allocation.

Metric	Description
* Mem|Contention (%)	This metric is used to monitor ESXi memory usage. When the value is high, it means the ESXi is using a good percentage of available memory. You may need to add more memory to other memory-related metrics. Key: mem|host_contentionPct
Mem|Contention (KB)	Host contention in kilobytes. Key: mem|host_contention
Mem|Host Usage (KB)	Machine usage in kilobytes. Key: mem|host_usage
Mem|Machine Demand (KB)	Host demand in kilobytes. Key: mem|host_demand
Mem|Overall Memory used to run VMs on Host (KB)	Overall memory used to run virtual machines on the host in kilobytes. Key: mem|host_usageVM
Mem|Provisioned Memory (KB)	Provisioned memory in kilobytes. Key: mem|host_provisioned
Mem|Minimum Free Memory (KB)	Minimum free memory. Key: mem|host_minfree
Mem|Reserved Capacity (%)	Percent reserved capacity. Key: mem|reservedCapacityPct
Mem|Usable Memory (KB)	Usable memory in kilobytes. Key: mem|host_usable
* Mem|Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem|host_usagePct
Mem|ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem|host_systemUsage
Mem|Guest Active (KB)	Amount of memory that is actively used. Key: mem|active_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory. Key: mem|consumed_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Heap (KB)	Amount of memory allocated for heap. Key: mem|heap_average
Mem|Heap Free (KB)	Amount of free space in the heap. Key: mem|heapfree_average
Mem|VM Overhead (KB)	Memory overhead reported by host. Key: mem|overhead_average

Metric	Description
Mem|Reserved Capacity (KB)	Reserved capacity in kilobytes. Key: mem|reservedCapacity_average
Mem|Shared (KB)	Amount of shared memory in kilobytes. Key: mem|shared_average
Mem|Shared Common (KB)	Amount of shared common memory in kilobytes. Key: mem|sharedcommon_average
Mem|Swap In (KB)	Amount of memory swapped in. Key: mem|swapin_average
Mem|Swap Out KB)	Amount of memory swapped out. Key: mem|swapout_average
Mem|Swap Used (KB)	Amount of memory used for swapped space in kilobytes. Key: mem|swapused_average
Mem|VM kernel Usage (KB)	Amount of memory used by the VM kernel. Key: mem|sysUsage_average
Mem|Unreserved (KB)	Amount of unreserved memory in kilobytes. Key: mem|unreserved_average
* Mem|Balloon (KB)	This metric shows the total amount of memory currently used by the VM memory control. This memory was reclaimed from the respective VMs at some point in the past, and was not returned. Use this metric to monitor how much VM memory has been reclaimed by ESXi through memory ballooning. The presence of ballooning indicates the ESXi has been under memory pressure. The ESXi activates ballooning when consumed memory reaches a certain threshold. Look for increasing size of ballooning. This indicates that there has been a shortage of memory more than once. Look for size fluctuations which indicate the ballooned out page was actually required by the VM. This translates into a memory performance problem for the VM requesting the page, since the page must first be brought back from the disk. Key: mem|vmmemctl_average
Mem|Zero (KB)	Amount of memory that is all zero. Key: mem|zero_average
Mem|State (0-3)	Overall state of the memory. The value is an integer between 0 (high) and 3 (low). Key: mem|state_latest
Mem|Usage (KB)	Host memory use in kilobytes. Key: mem|host_usage

Metric	Description
Mem|Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem|usage_average
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval in kilobyte per second. Key: mem|swapinRate_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory to disk during the current interval in kilobytes per second. Key: mem|swapoutRate_average
Mem|Active Write (KB)	Average active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Average memory compression in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Average compression rate in kilobytes per second. Key: mem|compressionRate_average
Mem|Decompression Rate (KBps)	Decompression rate in kilobytes per second. Key: mem|decompressionRate_average
Mem|Total Capacity (KB)	Total capacity in kilobytes. Key: mem|totalCapacity_average
Mem|Latency	Percentage of time the VM is waiting to access swapped or compressed memory. Key: mem|latency_average
Mem|Capacity Contention	Capacity Contention. Key: mem|capacity.contention_average
Mem|Swap In Rate from Host Cache	Rate at which memory is being swapped from host cache into active memory. Key: mem|llSwapInRate_average
Mem|Swap In from Host Cache	Amount of memory swapped-in from host cache. Key: mem|llSwapIn_average
Mem|Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem|llSwapOutRate_average
Mem|Swap Out to Host Cache	Amount of memory swapped-out to host cache. Key: mem|llSwapOut_average
Mem|Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem|llSwapUsed_average

Metric	Description
Mem|Low Free Threshold	Threshold of free host physical memory below which ESX will begin reclaiming memory from VMs through ballooning and swapping. Key: mem|lowfreethreshold_average
Mem|VM Memory Workload Disparity	Percentage Memory workload disparity among the VMs on the Host. Key: mem|vmWorkloadDisparityPct
Mem|Active Host Load For Balance (Long Term)	Active Host Load For Balance (Long Term). Key: mem|active_longterm_load
Mem|Active Host Load For Balance (Short Term)	Active Host Load For Balance (Short Term). Key: mem|active_shortterm_load

Network Metrics for Host Systems

Network metrics provide information about network performance.

Metric	Description
* Net|Packets Received per second	This metric shows the number of packets received during the collection interval. Use this metric to monitor the network usage of the ESXi. Key: net|packetsRxPerSec
* Net|Packets Transmitted per second	This metric shows the number of packets transmitted during the collection interval. Key: net|packetsTxPerSec
Net|Packets per second	Number of packets transmitted and received per second. Key: net|packetsPerSec
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|I/O Usage Capacity	I/O Usage Capacity. Key: net|usage_capacity
Net|Max Observed Throughput	Max observed rate of network throughput. Key: net|maxObserved_KBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps
Net|Demand (%)	Percent demand. Key: net|demand
Net|Data Transmit Rate (KBps)	Average amount of data transmitted per second. Key: net|transmitted_average
Net|Data Receive Rate (KBps)	Average amount of data received per second. Key: net|received_average

Metric	Description
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
Net|Received Packets Dropped	Number of received packets dropped in the performance interval. Key: net|droppedRx_summation
Net|Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval. Key: net|droppedTx_summation
* Net|Packets Dropped (%)	This metric shows the percentage of received and transmitted packets dropped during the collection interval. This metric is used to monitor reliability and performance of the ESXi network. When a high value is displayed, the network is not reliable and performance suffers. Key: net|droppedPct
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|bytes Rx (KBps)	Average amount of data received per second. Key: net|bytesRx_average
Net|bytes Tx (KBps)	Average amount of data transmitted per second. Key: net|bytesTx_average
Net|Broadcast Packets Received	Number of broadcast packets received during the sampling interval. Key: net|broadcastRx_summation
Net|Broadcast Packets Transmitted	Number of broadcast packets transmitted during the sampling interval. Key: net|broadcastTx_summation
Net|Error Packets Received	Number of packets with errors received. Key: net|errorsRx_summation
Net|Error Packets Transmitted	Number of packets with errors transmitted. Key: net|errorsTx_summation
Net|Multicast Packets Received	Number of multicast packets received. Key: net|multicastRx_summation
Net|Multicast Packets Transmitted	Number of multicast packets transmitted. Key: net|multicastTx_summation
Net|FT Throughput Usage	FT Throughput Usage. Key: net|throughput.usage.ft_average
Net|HBR Throughput Usage	HBR Throughput Usage. Key: net|throughput.usage.hbr_average

Metric	Description
Net|iSCSI Throughput Usage	iSCSI Throughput Usage. Key: net|throughput.usage.iscsi_average
Net|NFS Throughput Usage	NFS Throughput Usage. Key: net|throughput.usage.nfs_average
Net|VM Throughput Usage	VM Throughput Usage. Key: net|throughput.usage.vm_average
Net|vMotion Throughput Usage	vMotion Throughput Usage. Key: net|throughput.usage.vmotion_average
Net|Unknown Protocol Frames Received	Number of frames with unknown protocol received. Key: net|unknownProtos_summation

System Metrics for Host Systems

System metrics provide information about the amount of CPU that resources and other applications use.

Metric	Description
Sys|Power On	1 if the host system is powered on, 0 if the host system is powered off, or -1 if the power state is unknown. Key: sys|poweredOn
Sys|Uptime (seconds)	Number of seconds since the last system startup. Key: sys|uptime_latest
Sys|Disk Usage (%)	Percent disk use. Key: sys|diskUsage_latest
Sys|Resource CPU Usage (MHz)	Amount of CPU that the Service Console and other applications use. Key: sys|resourceCpuUsage_average
Sys|Resource CPU Active (1 min. average)	Percentage of resource CPU that is active. Average value during a one-minute period. Key: sys|resourceCpuAct1_latest
Sys|Resource CPU Active (%) (5 min. average)	Percentage of resource CPU that is active. Average value during a five-minute period. Key: sys|resourceCpuAct5_latest
Sys|Resource CPU Alloc Max (MHz)	Maximum resource CPU allocation in megahertz. Key: sys|resourceCpuAllocMax_latest
Sys|Resource CPU Alloc Min (MHz)	Minimum resource CPU allocation in megahertz. Key: sys|resourceCpuAllocMin_latest
Sys|Resource CPU Alloc Shares	Number of resource CPU allocation shares. Key: sys|resourceCpuAllocShares_latest
Sys|Resource CPU Max Limited (%) (1 min. average)	Percent of resource CPU that is limited to the maximum amount. Average value during a one-minute period. Key: sys|resourceCpuMaxLimited1_latest

Metric	Description
Sys|Resource CPU Max Limited (%) (5 min. average)	Percentage of resource CPU that is limited to the maximum amount. Average value during a five-minute period. Key: sys|resourceCpuMaxLimited5_latest
Sys|Resource CPU Run1 (%)	Percent resource CPU for Run1. Key: sys|resourceCpuRun1_latest
Sys|Resource CPU Run5 (%)	Percent resource CPU for Run5. Key: sys|resourceCpuRun5_latest
Sys|Resource Memory Alloc Max (KB)	Maximum resource memory allocation in kilobytes. Key: sys|resourceMemAllocMax_latest
Sys|Resource Memory Alloc Min (KB)	Minimum resource memory allocation in kilobytes. Key: sys|resourceMemAllocMin_latest
Sys|Resource Memory Alloc Shares	Number of resource memory shares allocated. Key: sys|resourceMemAllocShares_latest
Sys|Resource Memory Cow (KB)	Cow resource memory in kilobytes. Key: Sys|resourceMemCow_latest
Sys|Resource Memory Mapped (KB)	Mapped resource memory in kilobytes. Key: ys|resourceMemMapped_latest
Sys|Resource Memory Overhead (KB)	Resource memory overhead in kilobytes. Key: sys|resourceMemOverhead_latest
Sys|Resource Memory Shared (KB)	Shared resource memory in kilobytes. Key: sys|resourceMemShared_latest
Sys|Resource Memory Swapped (KB)	Swapped resource memory in kilobytes. Key: sys|resourceMemSwapped_latest
Sys|Resource Memory Touched (KB)	Touched resource memory in kilobytes. Key: sys|resourceMemTouched_latest
Sys|Resource Memory Zero (KB)	Zero resource memory in kilobytes. Key: sys|resourceMemZero_latest
Sys|Resource Memory Consumed	Resource Memory Consumed Latest (KB). Key: sys|resourceMemConsumed_latest
Sys|Resource File descriptors usage	Resource File descriptors usage (KB). Key: sys|resourceFdUsage_latest
Sys|vMotion Enabled	1 if vMotion is enabled or 0 if vMotion is not enabled. Key: sys|vmotionEnabled
Sys|Not in Maintenance	Not in maintenance. Key: sys|notInMaintenance

Management Agent Metrics for Host Systems

Management agent metrics provide information about memory use.

Metric	Description
Management Agent|Memory Used (%)	Amount of total configured memory that is available for use. Key: managementAgent|memUsed_average
Management Agent|Memory Swap Used (KB)	Sum of the memory swapped by all powered-on virtual machines on the host. Key: managementAgent|swapUsed_average
Management Agent|Memory Swap In (KBps)	Amount of memory that is swapped in for the Service Console. Key: managementAgent|swapIn_average
Management Agent|Memory Swap Out (KBps)	Amount of memory that is swapped out for the Service Console. Key: managementAgent|swapOut_average
Management Agent|CPU Usage	CPU usage. Key: managementAgent|cpuUsage_average

Storage Path Metrics for Host Systems

Storage path metrics provide information about data storage use.

Metric	Description
StoragePath|Total Latency (ms)	Total latency in milliseconds. Key: storagePath|totalLatency
StoragePath|Total Usage (KBps)	Total latency in kilobytes per second. Key: storagePath|usage
StoragePath|Read Rate (KBps)	Rate of reading data from the virtual disk. Key: storagePath|read_average
StoragePath|Write Rate (KBps)	Rate of writing data. Key: storagePath|write_average
StoragePath|Commands per second	Average number of commands issued per second during the collection interval. Key: storagePath|commandsAveraged_average
StoragePath|Reads per second	Average number of read commands issued per second during the collection interval. Key: storagePath|numberReadAveraged_average
StoragePath|Writes per second	Average number of write commands issued per second during the collection interval. Key: storagePath|totalWriteLatency_average
StoragePath|Writes per second	Average number of write commands issued per second during the collection interval. Key: storagePath|numberWriteAveraged_average
StoragePath|Read Latency (ms)	Average amount of time for a read operation by the storage adapter. Key: storagePath|totalReadLatency_average

Metric	Description
StoragePath|Highest Latency	Highest Latency. Key: storagePath|maxTotalLatency_latest
StoragePath|Storage Path Name	Storage path name. Key: storagePath|storagePathName

Storage Adapter Metrics for Host Systems

Storage adapter metrics provide information about data storage use.

Metric	Description
Storage Adapter|Total Usage (KBps)	Total latency. Key: storageAdapter|usage
Storage Adapter|Port WWN	Port World Wide Name. Key: storageAdapter|portWWN
Storage Adapter|Commands per second	Average number of commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|commandsAveraged_average
Storage Adapter|Reads per second	Average number of read commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|numberReadAveraged_average
Storage Adapter|Writes per second	Average number of write commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|numberWriteAveraged_average
Storage Adapter|Read Rate (KBps)	Rate of reading data by the storage adapter. Key: storageAdapter|read_average
* Storage Adapter|Read Latency (ms)	This metric shows the average amount of time for a read operation by the storage adapter. Use this metric to monitor the storage adapter read operation performance. A high value means that the ESXi is performing a slow storage read operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter|totalReadLatency_average
* Storage Adapter|Write Latency (ms)	This metric shows the average amount of time for a write operation by the storage adapter. Use this metric to monitor the storage adapter write performance operation. A high value means that the ESXi is performing a slow storage write operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter|totalWriteLatency_average
Storage Adapter|Write Rate (KBps)	Rate of writing data by the storage adapter. Key: storageAdapter|write_average
Storage Adapter|Demand	Demand. Key: storageAdapter|demand

Metric	Description
Storage Adapter|Highest Latency	Highest Latency. Key: torageAdapter|maxTotalLatency_latest
Storage Adapter|Outstanding Requests	Outstanding Requests. Key: storageAdapter|outstandingIOs_average
Storage Adapter|Queue Depth	Queue Depth. Key: storageAdapter|queueDepth_average
Storage Adapter|Queue Command Latency (ms)	The average time spent in the ESX Server VM Kernel queue per command. Key: storageAdapter|queueLatency_average
Storage Adapter|Queued	Queued. Key: storageAdapter|queued_average

Storage Metrics for Host Systems

Storage metrics provide information about storage use.

Metric	Description
Storage|Commands per second	Average number of commands issued per second during the collection interval. Key: storage|commandsAveraged_average
Storage|Read Latency (ms)	Average amount of time for a read operation in milliseconds. Key: storage|totalReadLatency_average
Storage|Read Rate (KBps)	Read throughput rate in kilobytes. Key: storage|read_average
Storage|Reads per second	Average number of read commands issued per second during the collection interval. Key: storage|numberReadAveraged_average
Storage|Total Latency (ms)	Total latency in milliseconds. Key: storage|totalLatency_average
Storage|Total Usage (KBps)	Total throughput rate in kilobytes per second. Key: storage|usage_average
Storage|Write Latency (ms)	Average amount of time for a write operation in milliseconds. Key: storage|totalWriteLatency_average
Storage|Write Rate (KBps)	Write throughput rate in kilobytes per second. Key: storage|write_average
Storage|Writes per second	Average number of write commands issued per second during the collection interval. Key: storage|numberWriteAveraged_average

Sensor Metrics for Host Systems

Sensor metrics provide information about host system cooling.

Metric	Description
Sensor|Fan|Speed (%)	Percent fan speed. Key: Sensor|fan|currentValue
Sensor|Fan|Health State	Fan health state. Key: Sensor|fan|healthState
Sensor|Temperature|Temp C	Fan temperature in centigrade. Key: Sensor|temperature|currentValue
Sensor|Temperature|Health State	Fan health state. Key: Sensor|temperature|healthState

Power Metrics for Host Systems

Power metrics provide information about host system power use.

Metric	Description
Power|Energy (Joule)	Host power use in joules. Key: power|energy_summation
Power|Power (Watt)	Host power use in watts. Key: power|power_average
Power|Power Cap (Watt)	Host power capacity in watts. Key: power|powerCap_average

Disk Space Metrics for Host Systems

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared disk space in gigabytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disks	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Shared Used (GB)	Used shared disk space in gigabytes. Key: diskspace|shared
Diskspace|Snapshot	Disk space used by snapshots in gigabytes. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Disk space used by virtual disks in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Disk space used by virtual machines in gigabytes. Key: diskspace|used
Diskspace|tTotal disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage

Metric	Description
Diskspace|Total disk spacey	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned .

Summary Metrics for Host Systems

Summary metrics provide information about overall host system performance.

Metric	Description
* Summary|Number of Running VMs	This metric shows the number of VMs running on the host during the last metric collection time. Large spikes of running VMs might be a reason for CPU or memory spikes as more resources will be used in the host. Number of Running VMs gives you a good indicator of how many requests the ESXi host must juggle. This excludes powered off VMs as they do not impact ESXi performance. A change in this number in your environment can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, as all the VMs will go down (or be relocated by HA) if the ESXi crashes. Look for any correlation between spikes in the number of running VMs and spikes in other metrics such as CPU Contention/Memory Contention. Key: summary|number_running_vms
Summary|Maximum Number of VMs	Maximum number of virtual machines Key: summary|max_number_vms
* Summary|Number of vMotions	This metric shows the number of vMotions that occurred in the host in the last X minutes. The number of vMotions is a good indicator of stability. In a healthy environment, this number should be stable and relatively low. Look for correlation between vMotions and spikes in other metrics such as CPU/Memory contention. The vMotion should not create any spikes, however, the VMs moved into the host might create spikes in memory usage, contention and CPU demand and contention. Key: summary|number_vmotion
Summary|Total Number of Datastores	Total Number of Datastores. Key: summary|total_number_datastores
Summary|Number of VCPUs on Powered On VMs	Total number of VCPUs of Virtual Machines that are powered on. Key: summary|number_running_vcpus

Metric	Description
Summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
Summary|Workload Indicator (%)	Percent workload indicator. Key: summary|workload_indicator

HBR Metrics for Host Systems

Host-based replication (HBR) metrics provide information about vSphere replication.

Metric	Description
HBR|Replication Data Received Rate	Replication Data Received Rate. Key: hbr|hbrNetRx_average
HBR|Replication Data Transmitted Rate	Replication Data Transmitted Rate. Key: hbr|hbrNetTx_average
HBR|Replicated VM Count	Number of replicated virtual machines. Key: hbr|hbrNumVms_average

Cluster Compute Resource Metrics

vRealize Operations Manager collects configuration, storage, disk space, CPU use, disk, memory, network, power, and summary metrics for cluster compute resources.

Cluster Compute Resource metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the clusters in your environment.

Configuration Metrics for Cluster Compute Resources

Configuration metrics provide information about configuration settings.

Metric	Description
Configuration|Failover Level	DAS configuration failover level. Key: configuration|dasconfig|failoverLevel
Configuration|Active Admission Control Policy	DAS configuration active admission control policy. Key: configuration|dasconfig| activeAdministrationControlPolicy

Metric	Description
Configuration|CPU Failover Resources Percent	Percent CPU failover resources for DAS configuration admission control policy. Key: configuration|dasconfig|adminissionControlPolicy|cpuFailoverResourcesPercent
Configuration|Memory Failover Resources Percent	Percent memory failover resources for DAS configuration admission control policy. Key: configuration|dasconfig|adminissionControlPolicy|memoryFailoverResourcesPercent

Storage Metrics for Cluster Compute Resources

Storage metrics provide information about storage use.

Metric	Description
Storage|Total Usage	Total throughput rate in kilobytes per second. Key: storage|usage_average

Disk Space Metrics for Cluster Compute Resources

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Virtual machine used (GB)	Space used by virtual machine files in gigabytes. Key: diskspace|used
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned
Diskspace|Virtual Disk Used (GB)	Space used by virtual disks in gigabytes. Key: diskspace|diskused
Diskspace|Snapshot Space (GB)	Space used by snapshots in gigabytes. Key: diskspace|snapshot
Diskspace|Shared Used (GB)	Shared used space in gigabytes. Key: diskspace|shared
Diskspace|Not Shared (GB)	Space used by VMs that is not shared. Key: diskspace|notshared

CPU Usage Metrics for Cluster Compute Resources

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|Capacity Usage	This metric shows the percentage of the capacity used. Key: cpu|capacity_usagepct_average
* CPU|CPU Contention (%)	This metric is an indicator of the overall contention for CPU resources that occurs across the workloads in the cluster. When contention occurs, it means that some of the virtual machines are not immediately getting the CPU resources they are requesting. Use this metric to identify when a lack of CPU resources might be causing performance issues in the cluster. This metric is the sum of the CPU contention across all hosts in the cluster averaged over two times the number of physical CPUs in the cluster to account for hyper-threading. CPU contention takes into account: ■ CPU Ready ■ CPU Co-stop ■ Power management ■ Hyper threading This metric is more accurate than CPU Ready since it takes into account CPU Co-stop and Hyper threading. When using this metric, the number should be lower than the performance you expect. If you expect performance at 10%, then the number should be lower than 10%. Since this value is averaged across all hosts in the cluster, you may find that some hosts have a higher CPU contention while others are lower. In order to ensure that vSphere spreads out the running workloads across hosts, consider enabling a fully-automated DRS in the cluster. Key: cpu|capacity_contentionPct
* CPU|Demand (%)	This metric is an indicator of the overall demand for CPU resources by the workloads in the cluster. It shows the percentage of CPU resources that all the virtual machines would use if there were no CPU contention or CPU limits set. It represents the average active CPU load in the past five minutes. Key: cpu|demandPct
CPU|Demand (MHz)	Demand in megahertz. Key: cpu|demandmhz
CPU|IO Wait	IO wait time in milliseconds. Key: cpu|iowait
CPU|Number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Host Provisioned Capacity	Provisioned CPU capacity in megahertz. Key: cpu|capacity_provisioned
CPU|Provisioned vCPUs	Number of provisioned CPU cores. Key: cpu|corecount_provisioned

Metric	Description
CPU|Reserved Capacity	The sum of the reservation properties of the (immediate) children of the host's root resource pool in megahertz. Key: cpu|reservedCapacity_average
CPU|Wait	CPU time spent on idle state in milliseconds. Key: cpu|wait
CPU|Usage (MHz)	Average CPU use in megahertz. Key: cpu|usagemhz_average
CPU|Total Capacity	Total CPU capacity in megahertz. Key: cpu|totalCapacity _average
CPU|Demand	CPU Demand. Key: cpu|demand_average
CPU|Overhead	Amount of CPU overhead. Key: cpu|overhead_average
CPU|Demand without overhead	Value of demand excluding any overhead. Key: cpu|demand_without_overhead
CPU|Provisioned Capacity	Provisioned Capacity (MHz). Key: cpu|vm_capacity_provisioned
CPU|Number of hosts stressed	Number of hosts stressed. Key: cpu|num_hosts_stressed
CPU|Stress Balance Factor	Stress Balance Factor. Key: cpu|stress_balance_factor
CPU|Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: cpu|min_host_capacity_remaining
CPU|Workload Balance Factor	Workload Balance Factor. Key: cpu|workload_balance_factor
CPU|Highest Provider Workload	Highest Provider Workload. Key: cpu|max_host_workload
CPU|Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: cpu|host_workload_disparity
CPU|Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: cpu|host_stress_disparity

Disk Metrics for Cluster Compute Resources

Disk metrics provide information about disk use.

Metric	Description
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Command Latency and Physical Device Command Latency metrics. Key: disk|totalLatency_average
Disk|Disk Read Latency	Average amount of time for a read operation from the virtual disk. The total latency is the sum of Kernel latency and device latency. Key: disk|totalReadLatency_average
Disk|Disk Write Latency	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalWriteLatency_averag
Disk|Read Rate (KBps)	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_averag
Disk|Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. Key: disk|usage_average
Disk|Write Rate (KBps)	Number of times data was written to disk during the collection interval. Key: disk|numberWrite_summation
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average
Disk|Read Requests	Amount of data read from the disk during the collection interval. Key: disk|read_average
Disk|Write Requests	Amount of data written to the disk during the collection interval. Key: disk|write_average
Disk|Commands Issued	Number of disk commands issued during the collection interval. Key: disk|commands_summation
Disk|Total Queued Outstanding operations	Sum of queued operation and outstanding operations. Key: disk|sum_queued_oio
Disk|Max Observed OIO	Max observed outstanding IO for a disk. Key: disk|max_observed

Memory Metrics for Cluster Compute Resources

Memory metrics provide information about memory use and allocation.

Metric	Description
Mem|Active Write (KB)	Active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Average compression in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Average compression rate in kilobytes. Key: mem|compressionRate_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory. Key: mem|consumed_average
* Mem|Contention (%)	This metric is an indicator of the overall contention for memory resources that occurs across the workloads in the cluster. When contention occurs, it means that some portion of the VMs are not immediately getting the memory resources that they are requesting. Use this metric to identify when lack of memory resources might be causing performance issues in the cluster. Key: mem|host_contentionPct
Mem|Contention (KB)	Contention in kilobytes. Key: mem|host_contention
Mem|Decompression Rate (KBps	Decompression rate in kilobytes. Key: mem|decompressionRate_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Guest Active (KB)	Amount of memory that is actively used. Key: mem|active_average
Mem|Heap (KB)	Amount of memory allocated for heap. Key: mem|heap_average
Mem|Heap Free (KB)	Free space in the heap. Key: mem|heapfree_average
* Mem|Balloon	This metric shows the amount of memory currently used by the virtual machine memory control. It is only defined at the VM level. Key: mem|vmmemctl_average
Mem|VM Overhead (KB)	Memory overhead reported by host. Key: mem|overhead_average
Mem|Provisioned Memory (KB)	Provisioned memory in kilobytes. Key: mem|host_provisioned
Mem|Reserved Capacity (KB)	Reserved capacity in kilobytes. Key: mem|reservedCapacity_average

Metric	Description
Mem|Shared (KB)	Amount of shared memory. Key: mem|shared_average
Mem|Shared Common (KB)	Amount of shared common memory. Key: mem|sharedcommon_average
Mem|Swap In (KB)	Amount of memory that is swapped in for the service console. Key: mem|swapin_average
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem|swapinRate_average
Mem|Swap Out (KB)	Amount of memory that is swapped out for the service console. Key: mem|swapout_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory into disk during the current interval. Key: mem|swapoutRate_average
Mem|Swap Used (KB)	Amount of memory used for swap space. Key: mem|swapused_average
Mem|Total Capacity (KB)	Total capacity in kilobytes. Key: mem|totalCapacity_average
Mem|Reserved (KB)	Amount of unreserved memory. Key: mem|unreserved_average
Mem|Usable Memory (KB)	Usable memory in kilobytes. Key: mem|host_usable
Mem|Usage/Usable	Percent memory used. Key: mem|host_usagePct
Mem|Host Usage (KB)	Memory use in kilobytes. Key: mem|host_usage
Mem|Machine Demand	Memory Machine Demand in KB. Key: mem|host_demand
Mem|ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem|host_systemUsage
* Mem|Usage (%)	This metric shows the portion of the total memory in all hosts in the cluster that is being used. This metric is the sum of memory consumed across all hosts in the cluster divided by the sum of physical memory across all hosts in the cluster. ∑ memory consumed on all hosts ---------- X 100% ∑ physical memory on all hosts
Mem|Usage (KB)	Memory currently in use as a percentage of total available memory. Key: mem|usage_average

Metric	Description
Mem|VM kernel Usage (KB)	Amount of memory that the VM kernel uses. Key: mem|sysUsage_average
Mem|Zero (KB)	Amount of memory that is all 0. Key: mem|zero_average
Mem|Number of Hosts Stressed	Number of hosts stressed. Key: mem|num_hosts_stressed
Mem|Stress Balance Factor	Stress balance factor. Key: mem|stress_balance_factor
Mem|Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: mem|min_host_capacity_remaining
Mem|Workload Balance Factor	Workload balance factor. Key: mem|workload_balance_factor
Mem|Highest Provider Workload	Highest provider workload. Key: mem|max_host_workload
Mem|Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: mem|host_workload_disparity
Mem|Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: mem|host_stress_disparity

Network Metrics for Cluster Compute Resources

Network metrics provide information about network performance.

Metric	Description
Net|Data Receive Rate (KBps)	Average amount of data received per second. Key: net|received_average
Net|Data Transmit Rate (KBps)	Average amount of data transmitted per second. Key: net|transmitted_average
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|Packets Dropped (%)	Percentage of packets dropped. Key: net|droppedPct
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
Net|Received Packets Dropped	Number of received packets dropped in the performance interval. Key: net|droppedRx_summation
Net|Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval. Key: net|droppedTx_summation

Metric	Description
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|Max Observed Throughput	Max observed rate of network throughput. Key: net|maxObservedKBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps

Datastore Metrics for Cluster Compute Resources

Datastore metrics provide information about Datastore use.

Metric	Description
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average
Datastore|Read Rate	Amount of data read in the performance interval. Key: datastore|read_average
Datastore|Write Rate	Amount of data written to disk in the performance interval. Key: datastore|write_average

Cluster Services Metrics for Cluster Compute Resources

Cluster Services metrics provide information about cluster services.

Metric	Description
ClusterServices|Effective CPU Resources (MHz)	VMware DRS effective CPU resources available. Key: clusterServices|effectivecpu_average
ClusterServices|Effective Memory Resources (KB)	VMware DRS effective memory resources available. Key: clusterServices|effectivemem_average

Power Metrics for Cluster Compute Resources

Power metrics provide information about power use.

Metric	Description
Power|Energy (Joule)	Energy use in joules. Key: power|energy_summation
Power|Power (Watt)	Average power use in watts. Key: power|power_average
Power|Power Cap (Watt)	Average power capacity in watts. Key: power|powerCap_average

Summary Metrics for Cluster Compute Resources

Summary metrics provide information about overall performance.

Metric	Description
Summary|Number of Running Hosts	Number of running hosts. Key: summary|number_running_hosts
* Summary|Number of Running VMs	This metric shows the total number of VMs running on all hosts in the cluster. Key: summary|number_running_vms
* Summary|Number of vMotions	This metric shows the number of vMotions that occurred during the last collection cycle. When using this metric, look for a low number which indicates that the cluster is able to serve its VMs. A vMotion can impact VM performance during the stun time. Key: summary|number_vmotion
Summary|Total Number of Hosts	Total number of hosts. Key: summary|total_number_hosts
Summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
Summary|Maximum Number of VMs	Maximum Number of virtual machines. Key: summary|max_number_vms

Metric	Description
Summary|Workload Indicator	Percent workload indicator. Key: summary|workload_indicator
Summary|Total Number of Datastores	Total number of datastores. Key: summary|total_number_datastores
Summary|Number of VCPUs on Powered On VMs	Number of virtual CPUs on powered-on virtual machines. Key: summary|number_running_vcpus
Summary|Average Running VM Count per Running Host	Average number of running virtual machines per running host. Key: summary|avg_vm_density
Summary|Average Provisioned Capacity (MHz) per Running VM	Average provisioned capacity, in megahertz, per running virtual machine. Key: summary|avg_vm_cpu
Summary|Average Provisioned Memory (KB) per Running VM	Average provisioned memory, in kilobytes. per running virtual machine. Key: summary|avg_vm_mem

Resource Pool Metrics

vRealize Operations Manager collects configuration, CPU usage, memory, and summary metrics for resource pool objects.

Resource Pool metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

Configuration Metrics for Resource Pools

Configuration metrics provide information about memory and CPU allocation configuration.

Table 7-14. Configuration Metrics for Resource Pools

Metric Key	Metric Name	Description
config|mem_alloc_reservation	Memory Allocation Reservation	Memory Allocation Reservation.

CPU Usage Metrics for Resource Pools

CPU usage metrics provide information about CPU use.

Table 7-15. CPU Usage Metrics for Resource Pools

Metric Key	Metric Name	Description
cpu|capacity_demandEntitlementPct	Capacity Demand Entitlement (%)	CPU Capacity Demand Entitlement Percentage.
cpu|capacity_entitlement	Capacity entitlement (MHz)	CPU Capacity Entitlement.
cpu|capacity_contentionPct	CPU Contention (%)	CPU capacity contention.

Table 7-15. CPU Usage Metrics for Resource Pools (continued)

Metric Key	Metric Name	Description
cpu|demandmhz	Demand (MHz)	CPU demand in megahertz.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|usagemhz_average	Usage	Average CPU use in megahertz.
cpu|effective_limit	Effective limit	CPU effective limit.
cpu|reservation_used	Reservation Used	CPU reservation used.
cpu|estimated_entitlement	Estimated entitlement	CPU estimated entitlement.
cpu|dynamic_entitlement	Dynamic entitlement	CPU dynamic entitlement.
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead

Memory Metrics for Resource Pools

Memory metrics provide information about memory use and allocation.

Table 7-16. Memory Metrics for Resource Pools

Metric Key	Metric Name	Description
mem|vmmemctl_average	Balloon (KB)	Amount of memory currently used by the virtual machine memory control.
mem|compressionRate_average	Compression Rate (KBps)	Compression rate in kilobytes per second.
mem|consumed_average	Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory.
mem|host_contentionPct	Contention (%)	Macine contention percentage.
mem|guest_usage	Guest usage	Guest memory entitlement.
mem|guest_demand	Guest demand	Guest memory entitlement.
mem|host_contention	Contention (KB)	Machine contention in kilobytes.
mem|decompressionRate_average	Decompression Rate (KBps)	Decompression rate in kilobytes per second.
mem|granted_average	Granted (KB)	Average of memory available for use.
mem|active_average	Guest Active (KB)	Amount of memory that is actively used.
mem|overhead_average	VM Overhead (KB)	Memory overhead reported by host.
mem|shared_average	Shared (KB)	Amount of shared memory.
mem|reservation_used	Reservation Used	Memory Reservation Used.
mem|dynamic_entitlement	Dynamic Entitlement	Memory Dynamic Entitlement.
mem|effective_limit	Effective Limit	Memory Effective Limit.

Table 7-16. Memory Metrics for Resource Pools (continued)

Metric Key	Metric Name	Description
mem|swapinRate_average	swapinRate_average	Rate at which memory is swapped from disk into active memory during the interval.
mem|swapoutRate_average	swapoutRate_average	Rate at which memory is being swapped from active memory to disk during the current interval.
mem|swapped_average	Swapped (KB)	Amount of unreserved memory.
mem|usage_average	Usage (%)	Memory currently in use as a percentage of total available memory.
mem|zero_average	Zero (KB)	Amount of memory that is all zero.
mem|zipped_latest	Zipped (KB)	Latest zipped memory in kilobytes.
mem|swapin_average	Swap In (KB)	Amount of memory swapped in kilobytes.
mem|swapout_average	Swap Out (KB)	Amount of memory swapped out in kilobytes.
mem|swapused_average	Swap Used (KB)	Amount of memory used for swap space in kilobytes.
mem|guest_provisioned	Guest Configured Memory (KB)	Guest configured memory in kilobytes.

Summary Metrics for Resource Pools

Summary metrics provide information about overall performance.

Table 7-17. Summary Metrics for Resource Pools

Metric Key	Metric Name	Description
summary|number_running_vms	Number of Running VMs	Number of running virtual machines.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|iowait	IO Wait (ms)	IO wait time in milliseconds.

Datacenter Metrics

vRealize Operations Manager collects CPU usage, disk, memory, network, storage, disk space, and summary metrics for datacenter objects.

Datacenter metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

CPU Usage Metrics for Datacenters

CPU usage metrics provide information about CPU use.

Table 7-18. CPU Usage Metrics for Datacenters

Metric Key	Metric Name	Description
cpu|capacity_usagepct_average	Capacity Usage (%)	Percent capacity used.
cpu|capacity_contentionPct	CPU Contention (%)	CPU capacity contention.
cpu|demandPct	Demand (%)	CPU demand percentage.
cpu|demandmhz	Demand	Demand in megahertz.
cpu|demand_average	Demand (MHz)	CPU Demand.
cpu|overhead_average	Overhead (KB)	Amount of CPU overhead.
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|wait	Total Wait	CPU time spent on idle state.
cpu|numpackages	Number of CPU Sockets	Number of CPU sockets.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|capacity_provisioned	Host Provisioned Capacity (MHz)	Host provisioned capacity in megahertz.
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s).
cpu|reservedCapacity_average	Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool.
cpu|usagemhz_average	Usage	Average CPU usage in megahertz.
cpu|iowait	IO Wait	IO wait time in milliseconds.
cpu|vm_capacity_provisioned	Provisioned Capacity	Provisioned Capacity.
cpu|stress_balance_factor	Stress Balance Factor	Stress Balance Factor.
cpu|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining.
cpu|workload_balance_factor	Workload Balance Factor	Workload Balance Factor.
cpu|max_host_workload	Highest Provider Workload	Highest Provider Workload.
cpu|host_workload_disparity	Host workload Max-Min Disparity	Difference of Max and Min host workload in the container.
cpu|host_stress_disparity	Host stress Max-Min Disparity	Difference of Max and Min host stress in the container.

Disk Metrics for Datacenters

Disk metrics provide information about disk use.

Table 7-19. Disk Metrics for Datacenters

Metric Key	Metric Name	Description
disk|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection interval.
disk|totalLatency_average	Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Disk Command Latency and Physical Device Command Latency metrics.
disk|usage_average	Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.
disk|sum_queued_oio	Total queued outstanding operations	Sum of queued operations and outstanding operations.
disk|max_observed	Max observed OIO	Max observed IO for a disk.

Memory Metrics for Datacenters

Memory metrics provide information about memory use and allocation.

Table 7-20. Memory Metrics for Datacenters

Metric Key	Metric Name	Description
mem|host_contentionPct	Contention (%)	Machine Contention Percentage.
mem|host_demand	Machine Demand (KB)	Memory machine demand in kilobytes.
mem|host_systemUsage	ESX System Usage	Memory usage by the VM kernel and ESX user-level services.
mem|host_provisioned	Provisioned Memory (KB)	Provisioned host memory in kilobytes.
mem|reservedCapacity_average	Reserved Capacity (KB)	Reserved memory capacity in kilobytes.
mem|host_usable	Usable Memory (KB)	Usable host memory in kilobytes.
mem|host_usage	Host Usage	Host memory use in kilobytes.
mem|host_usagePct	Usage/Usable (%)	Percent host memory used.
mem|overhead_average	VM Overhead	Memory overhead reported by host.
mem|stress_balance_factor	Stress Balance Factor	Stress Balance Factor.
mem|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining.
mem|workload_balance_factor	Workload Balance Factor	Workload Balance Factor.
mem|max_host_workload	Highest Provider Workload	Highest Provider Workload.
mem|host_workload_disparity	Host workload Max-Min Disparity	Difference of Max and Min host workload in the container.
mem|host_stress_disparity	Host stress Max-Min Disparity	Difference of Max and Min host stress in the container.

Network Metrics for Datacenters

Network metrics provide information about network performance.

Table 7-21. Network Metrics for Datacenters

Metric Key	Metric Name	Description
net|droppedPct	Packets Dropped	Percentage of packets dropped.
net|maxObservedKBps	Max Observed Throughput	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput	Max observed received rate of network throughput.
net|transmitted_average	Data Transmit Rate	Average amount of data transmitted per second.
net|received_average	Data Receive Rate	Average amount of data received per second.
net|usage_average	Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.

Storage Metrics for Datacenters

Storage metrics provide information about storage use.

Table 7-22. Storage Metrics for Datacenters

Metric Key	Metric Name	Description
storage|usage_average	Total Usage	Total throughput rate.

Datastore Metrics for Datacenters

Datastore metrics provide information about Datastore use.

Table 7-23. Datastore Metrics for Datacenters

Metric Key	Metric Name	Description
datastore|maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore|maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations.

Table 7-23. Datastore Metrics for Datacenters (continued)

Metric Key	Metric Name	Description
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore| numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore| numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.

Disk Space Metrics for Datacenters

Disk space metrics provide information about disk use.

Table 7-24. Disk Space Metrics for Datacenters

Metric Key	Metric Name	Description
diskspace|used	Virtual machine used	Used virtual machine disk space in gigabytes.
diskspace|total_usage	Total disk space used	Total disk space used on all datastores visible to this object.
diskspace|total_capacity	Total disk space	Total disk space on all datastores visible to this object.
diskspace|total_provisioned	Total provisioned disk space	Total provisioned disk space on all datastores visible to this object.
diskspace|notshared	Not Shared (GB)	Unshared disk space in gigabytes.
diskspace|shared	Shared Used (GB)	Shared disk space in gigabytes.
diskspace|snapshot	Snapshot Space (GB)	Snapshot disk space in gigabytes.
diskspace|diskused	Virtual Disk Used (GB)	Used virtual disk space in gigabytes.
diskspace|numvmdisk	Number of Virtual Disks	Number of Virtual Disks.

Summary Metrics for Datacenters

Summary metrics provide information about overall performance.

Table 7-25. Summary Metrics for Datacenters

Metric Key	Metric Name	Description
summary|number_running_hosts	Number of Running Hosts	Number of hosts that are ON.
summary|number_running_vms	Number of Running VMs	Number of running virtual machines.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.

Table 7-25. Summary Metrics for Datacenters (continued)

Metric Key	Metric Name	Description
summary|total_number_clusters	Total Number of Clusters	Total number of clusters.
summary|total_number_hosts	Total Number of Hosts	Total number of hosts.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|total_number_datastores	Total Number of Datastores	Total number of datastores.
summary|number_running_vcpus	Number of VCPUs on Powered On VMs	Total number of VCPUs of virtual machines that are powered on.
summary|workload_indicator	Workload Indicator	Workload indicator.
summary|avg_vm_density	Average Running VM Count per Running Host	Average number of running virtual machines per running host.

Custom Datacenter Metrics

vRealize Operations Manager collects CPU usage, memory, summary, network, and datastore metrics for custom datacenter objects.

Custom datacenter metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

CPU Usage Metrics for Custom Datacenters

CPU usage metrics provide information about CPU use.

Table 7-26. CPU Usage Metrics for Custom Datacenters

Metric Key	Metric Name	Description
cpu|capacity_provisioned	Host Provisioned Capacity	Host provisioned capacity (MHz).
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s).
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|num_hosts_stressed	Number of hosts stressed	Number of hosts stressed.
cpu|stress_balance_factor	Stress Balance Factor	Stress balance factor.
cpu|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest provider capacity remaining.
cpu|workload_balance_factor	Workload Balance Factor	Workload balance factor.
cpu|max_host_workload	Highest Provider Workload	Highest provider workload.
cpu|host_workload_disparity	Host workload Max-Min Disparity	Host workload max-min disparity.
cpu|host_stress_disparity	Host stress Max-Min Disparity	Difference of max and min host stress in the container.

Memory Metrics for Custom Datacenters

Memory metrics provide information about memory use.

Table 7-27. Memory Metrics for Custom Datacenters

Metric Key	Metric Name	Description
mem|host_usable	Usable Memory	Usable memory.
mem|host_demand	Machine Demand	Memory machine demand in KB.
mem|num_hosts_stressed	Number of hosts stressed	Number of hosts stressed.
mem|stress_balance_factor	Stress Balance Factor	Stress balance factor.
mem|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest provider capacity remaining.
mem|workload_balance_factor	Workload Balance Factor	Workload balance factor.
mem|max_host_workload	Highest Provider Workload	Highest provider workload.
mem|host_workload_disparity	Host workload Max-Min Disparity	Host workload max-min disparity.
mem|host_stress_disparity		Host stress max-min disparity.

Summary Metrics for Custom Datacenters

Summary metrics provide information about overall performance.

Table 7-28. Summary Metrics for Custom Datacenters

Metric Key	Metric Name	Description
summary|number_running_vms	Number of Running VMs	Number of virtual machines that are ON.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.
summary|status	Status	Status of datacenter.

Network Metrics for Custom Datacenters

Network metrics provide information about network performance.

Table 7-29. Network Metrics for Custom Datacenters

Metric Key	Metric Name	Description
net|usage_average	Usage Rate	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.
net|maxObserved_KBps	Max Observed Throughput	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput	Max observed received rate of network throughput.

Table 7-29. Network Metrics for Custom Datacenters (continued)

Metric Key	Metric Name	Description
net|transmitted_average	Data Transmit Rate	Average amount of data transmitted per second.
net|received_average	Data REceive Rate	Average amount of data received per second.

Datastore Metrics for Custom Datacenters

Datastore metrics provide information about datastore use.

Table 7-30. Datastore Metrics for Custom Datacenters

Metric Key	Metric Name	Description
datastore|maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore|maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Max observer number of outstanding IO operations.
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore| numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore| numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read rate	Amount of data read in the performance interval.
datastore|write_average	Write rate	Amount of data written to disk in the performance interval.

Storage Pod Metrics

vRealize Operations Manager collects datastore and disk space metrics for storage pod objects.

Storage Pod metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

Table 7-31. Datastore Metrics for Storage Pods

Metric Key	Metric Name	Description
datastore|numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore|numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.
datastore|usage_average	Usage Average	Usage Average.
datastore|totalReadLatency_average	Read Latency	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.
datastore|totalWriteLatency_average	Write Latency	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.
datastore|totalLatency_average	Disk Command Latency	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.
datastore|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection interval.

Table 7-32. Diskspace Metrics for Storage Pods

Metric Key	Metric Name	Description
diskspace|disktotal	Total used	Total space used.
diskspace|freespace	Freespace	Unused space available on datastore.
diskspace|capacity	Capacity	Total capacity of datastore.
diskspace|used	Virtual Machine used	Space used by virtual machine files.
diskspace|snapshot	Snapshot Space	Space used by snapshots.

VMware Distributed Virtual Switch Metrics

vRealize Operations Manager collects network and summary metrics for VMware distributed virtual switch objects.

VMware Distributed Virtual Switch metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

Table 7-33. Network Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
network|port_statistics|rx_bytes	Total Ingress Traffic	Total ingress traffic (KBps).
network|port_statistics|tx_bytes	Total Egress Traffic	Total egress traffic (KBps).
network|port_statistics|ucast_tx_pkts	Egress Unicast Packets per second	Egress unicast packets per second.
network|port_statistics|mcast_tx_pkts	Egress Multicast Packets per second	Egress multicast packets per second.
network|port_statistics|bcast_tx_pkts	Egress Broadcast Packets per second	Egress broadcast packets per second.
network|port_statistics|ucast_rx_pkts	Ingress Unicast Packets per second	Ingress unicast packets per second.
network|port_statistics|mcast_rx_pkts	Ingress Multicast Packets per second	Ingress multicast packets per second.
network|port_statistics|bcast_rx_pkts	Ingress Broadcast Packets per second	Ingress broadcast packets per second.
network|port_statistics|dropped_tx_pkts	Egress Dropped Packets per second	Egress dropped packets per second.
network|port_statistics|dropped_rx_pkts	Ingress Dropped Packets per second	Ingress dropped packets per second.
network|port_statistics|rx_pkts	Total Ingress Packets per second	Total ingress packets per second.
network|port_statistics|tx_pkts	Total Egress Packets per second	Total egress packets per second.
network|port_statistics|utilization	Utilization	Use (KBps).
network|port_statistics|dropped_pkts	Total Dropped Packets per second	Total dropped packets per second.
network|port_statistics|dropped_pkts_pct	Percentage of Dropped Packets	Percentage of dropped packets.
network|port_statistics|maxObserved_rx_bytes	Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps).
network|port_statistics|maxObserved_tx_bytes	Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps).
network|port_statistics|maxObserved_utilization	Max Observed Utilization (KBps)	Max observed utilization (KBps).

Table 7-34. Summary Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
summary|max_num_ports	Maximum Number of Ports	Maximum number of ports.
summary|used_num_ports	Used Number of Ports	Used number of ports.
summary|num_blocked_ports	Number of Blocked Ports	Number of blocked ports.

Table 7-35. Host Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
host|mtu_mismatch	MTU Mismatch	Maximum Transmission Unit (MTU) mismatch.
host|teaming_mismatch	Teaming Mismatch	Teaming mismatch.
host|mtu_unsupported	Unsupported MTU	Unsupported MTU.
host|vlans_unsupported	Unsupported VLANs	Unsupported VLANs.
host|config_outofsync	Config Out Of Sync	Config Out Of Sync.
host|attached_pnics	Number of Attached pNICs	Number of attached physical NICs.

Distributed Virtual Port Group Metrics

The vCenter Adapter instance collects network and summary metrics for distributed virtual port groups.

Distributed Virtual Port Group metrics include capacity and badge metrics. See definitions in:

- Capacity and Project-Based Metrics
- Badge Metrics

Table 7-36. Network Metrics for Distributed Virtual Port Groups

Metric Key	Metric Name	Description
network|port_statistics|rx_bytes	Ingress Traffic	Ingress traffic (KBps).
network|port_statistics|tx_bytes	Egress Traffic	Egress traffic (KBps).
network|port_statistics|ucast_tx_pkts	Egress Unicast Packets per second	Egress unicast packets per second.
network|port_statistics|mcast_tx_pkts	Egress Multicast Packets per second	Egress multicast packets per second.
network|port_statistics|bcast_tx_pkts	Egress Broadcast Packets per second	Egress broadcast packets per second.
network|port_statistics|ucast_rx_pkts	Ingress Unicast Packets per second	Ingress unicast packets per second.
network|port_statistics|mcast_rx_pkts	Ingress Multicast Packets per second	Ingress multicast packets per second.
network|port_statistics|bcast_rx_pkts	Ingress Broadcast Packets per second	Ingress broadcast packets per second.
network|port_statistics|dropped_tx_pkts	Egress Dropped Packets per second	Egress dropped packets per second.
network|port_statistics|dropped_rx_pkts	Ingress Dropped Packets per second	Ingress dropped packets per second.
network|port_statistics|rx_pkts	Total Ingress Packets per second	Total Ingress packets per second.
network|port_statistics|tx_pkts	Total Egress Packets per second	Total Egress packets per second.
network|port_statistics|utilization	Utilization	Utilization (KBps).

Table 7-36. Network Metrics for Distributed Virtual Port Groups (continued)

Metric Key	Metric Name	Description
network|port_statistics|dropped_pkts	Total Dropped Packets per second	Total dropped packets per second.
network|port_statistics|dropped_pkts_pct	Percentage of Dropped Packets	Percentage of dropped packets.
network|port_statistics|maxObserved_rx_bytes	Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps).
network|port_statistics|maxObserved_tx_bytes	Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps).
network|port_statistics|maxObserved_utilization	Max Observed Utilization (KBps)	Max observed utilization (KBps).

Table 7-37. Summary Metrics for Distributed Virtual Port Groups

Metric Key	Metric Name	Description
summary|max_num_ports	Maximum Number of Ports	Maximum number of ports.
summary|used_num_ports	Used Number of Ports	Used number of ports.
summary|num_blocked_ports	Number of Blocked Ports	Number of blocked ports.

Datastore Metrics

vRealize Operations Manager collects capacity, device, and summary metrics for datastore objects.

Capacity metrics can be calculated for datastore objects. See Capacity and Project-Based Metrics.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the datastores in your environment.

Capacity Metrics for Datastores

Capacity metrics provide information about datastore capacity.

Metric	Description
* Capacity|Available Space (GB)	This metric shows the amount of free space that a datastore has available. Use this metric to know how much storage space is unused on the datastore. Try to avoid having too little free disk space in order to accommodate unexpected storage growth on the datastore. The exact size of the datastore is based on company policy. Key: capacity|available_space
Capacity|Data Store Capacity Contention	Datastore capacity contention. Key: capacity|contention

Metric	Description
* Capacity|Provisioned (GB)	This metric shows the amount of storage that was allocated to the virtual machines. Use this metric to know how much storage space is currently being used on the datastore. Check the metric trend to identify spikes or abnormal growth. Key: capacity|provisioned
* Capacity|Total Capacity (GB)	This metric shows the overall size of the datastore. Use this metric to know the total capacity of the datastore. Typically the size of the datastore should not be too small. VMFS datastore size has grown over the years as virtualization matures and larger vitual machines are now onboard. Ensure the size can handle enough virtual machines to avoid datastore sprawl. A best practise is to use 5 TB for VMFS and more for vSAN. Key: capacity|total_capacity
Capacity|Used Space (GB)	This metric shows the amount of storage that is being used on the datastore. Key: capacity|used_space
Capacity|Workload (%)	Capacity workload. Key: capacity|workload
Capacity|Uncommitted Space (GB)	Uncommitted space in gigabytes. Key: capacity|uncommitted
Capacity|Total Provisioned Consumer Space	Total Provisioned Consumer Space. Key: capacity|consumer_provisioned
* Capacity|Used Space (%)	This metric shows the amount of storage that is being used on the datastore. Use this metric to know the percentage of storage space being used on the datastore. When using this metric, verify that you have at least 20% of free storage. Less than this, and you may experience problems when a snapshot is not deleted. If you have more than 50% free storage space, you are not utilizing your storage in the best possible way. Key: capacity|usedSpacePct

Device Metrics for Datastores

Device metrics provide information about device performance.

Metric	Description
Devices|Bus Resets	This metric shows the number of bus resets in the performance interval. Key: devices|busResets_summation
Devices|Commands Aborted	This metric shows the number of disk commands aborted in the performance interval. Key: devices|commandsAborted_summation

Metric	Description
Devices|Commands Issued	This metric shows the number of disk commands issued in the performance interval. Key: devices|commands_summation
Devices|Disk Command Latency (ms)	This metric shows the average time taken for a command from the perspective of a guest operating system. This metric is the sum of Kernel Disk Command Latency and Physical Device Command Latency metrics. Key: devices|totalLatency_average
Devices|Disk Read Latency (ms)	This metric shows the average time taken for a read from the perspective of a guest operating system. This metric is the sum of the Kernel Disk Read Latency and Physical Device Read Latency metrics. Key: devices|totalReadLatency_averag
Devices|Disk Write Latency (ms)	This metric shows the average amount of time for a write operation to the datastore. Total latency is the sum of kernel latency and device latency. Key: devices|totalWriteLatency_average
Devices|Kernel Disk Command Latency (ms)	Average time spent in ESX Server V. Kernel per command. Key: devices|kernelLatency_average
Devices|Kernel Disk Read Latency (ms)	Average time spent in ESX host VM Kernel per read. Key: devices|kernelReadLatency_average
Devices|Kernel Disk Write Latency (ms)	Average time spent in ESX Server VM Kernel per write. Key: devices|kernelWriteLatency_average
Devices|Number of Running Hosts	Number of running hosts that are powered on. Key: devices|number_running_hosts
Devices|Number of Running VMs	Number of running virtual machines that are powered on. Key: devices|number_running_vms
Devices|Physical Device Command Latency (ms)	Average time taken to complete a command from the physical device. Key: devices|deviceLatency_average
Devices|Physical Device Read Latency (ms)	Average time taken to complete a read from the physical device. Key: devices|deviceReadLatency_average
Devices|Queue Command Latency (ms)	Average time spent in the ESX Server VM Kernel queue per command. Key: devices|queueLatency_average
Devices|Queue Read Latency (ms)	Average time spent in the ESX Server VM Kernel queue per read. Key: devices|queueReadLatency_average
Devices|Queue Write Latency (ms)	Average time spent in the ESX Server VM Kernel queue per write. Key: devices|queueWriteLatency_average

Metric	Description
Devices|Read Rate (KBps)	Amount of data read in the performance interval. Key: devices|read_average
Devices|Read Requests	Number of times data was read from the disk in the defined interval. Key: devices|numberRead_summation
Devices|Reads per second	Average number of read commands issued per second to the datastore during the collection interval. Key: devices|numberReadAveraged_average
Devices|Usage Average (KBps)	Average use in kilobytes per second. Key: devices|usage_average
Devices|Write Rate (KBps)	Amount of data written to disk in the performance interval. Key: devices|write_average
Devices|Write Requests	Number of times data was written to the disk in the defined interval. Key: devices|numberWrite_summation
Devices|Writes per second	Average number of write commands issued per second to the datastore during the collection interval. Key: devices|numberWriteAveraged_average
Devices|Commands per second	Average number of commands issued per second during the collection interval. Key: devices|commandsAveraged_average
Devices|Physical Device Write Latency (ms)	Average time taken to complete a write from the physical disk. Key: devices|deviceWriteLatency_average

Datastore Metrics for Datastores

Datastore metrics provide information about datastore use.

Metric	Description
* Datastore|Disk Command Latency (ms)	This metric shows the adjusted read and write latency at the datastore level. Adjusted means that the latency is taking into account the number of IOs. If your IO is read-dominated, the combined value is influenced by the reads. This is the average of all the VMs running in the datastore. Because it is an average, some VMs logically experience higher latency that the value shown by this metric. To see the worst latency experienced by any VM, use the Maximum VM Disk Latency metric. Use this metric to see the performance of the datastore. It is one of two key performance indicators for a datastore, the other being the Max Read Latency. The combination of Maximum and Average gives better insight into how well the datastore is coping with the demand. The number should be lower than the performance you expect. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Average use in kilobytes per second. Key: datastore|usage_average
Datastore|Read Latency (ms	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
* Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average
Datastore|Demand	Demand. Key: datastore|demand
Datastore|Demand Indicator	Demand Indicator. Key: datastore|demand_indicator
Datastore|Max Observed Reads per Second	Maximum observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate (KBps)	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read

Metric	Description
* Datastore|Max Observed Read Latency (ms)	This metric displays the maximum observed average amount of time for a read operation from the datastore. Use this metric to troubleshoot when the overall datastore latency is higher than expected. Look for a number that matches the overall latency. Total latency = kernel latency + device latency. Key: datastore|maxObserved_ReadLatency
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate (KBps)	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Write Latency (ms)	This metric displays the maximum observed average amount of time for a write operation from the datastore. Use this metric to troubleshoot when the overall datastore latency is higher than expected. Look for a number that matches the overall latency. Total latency = kernel latency + device latency. Key: datastore|maxObserved_WriteLatency
Datastore|Max Observed Number of Outstanding IO Operations	Maximum observed number of outstanding IO operations. Key: datastore|maxObserved_OIO
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
* Datastore|Reads per second (IOPS)	This metric displays the average number of read commands issued per second during the collection interval. Use this metric when the total IOPS is higher than expected. Drill down to see if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system. Key: datastore|numberReadAveraged_average

Metric	Description
* Datastore|Writes per second (IOPS)	This metric displays the average number of write commands issued per second during the collection interval. Use this metric when the total IOPS is higher than expected. Drill down to see if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system. Key: datastore|numberWriteAveraged_average
Datastore|Read rate	This metric displays the amount of data read in the performance interval. Key: datastore|read_average
Datastore|Write rate	This metric displays the amount of data written to disk in the performance interval. Key: datastore|write_average

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.

Disk Space Metrics for Datastores

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared space in gigabytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disk	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Provisioned Space (GB)	Provisioned space in gigabytes. Key: diskspace|provisioned
Diskspace|Shared Used (GB)	Shared used space in gigabytes. Key: diskspace|shared

Metric	Description
* Diskspace|Snapshot Space (GB)	This metric shows the amount of space taken by snapshots on a given database. Use this metric to know how much storage space is being used by virtual machine snapshots on the datastore. Check that the snapshot is using 0 GB or minimal space. Anything over 1 GB should trigger a warning. The actual value depends on how IO intensive the virtual machines in the datastore are. Run a DT on them to detect anomaly. Clear the snapshot within 24 hours, preferably as soon as you have finished backing up, or patching. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Virtual disk used space in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Virtual machine used space in gigabytes. Key: diskspace|used
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned
Diskspace|Total used (GB)	Total used space in gigabytes. Key: diskspace|disktotal
Diskspace|Swap File Space (GB)	Swap file space in gigabytes. Key: diskspace|swap
Diskspace|Other VM Space (GB)	Other virtual machine space in gigabytes. Key: diskspace|otherused
Diskspace|Freespace (GB)	Unused space available on datastore. Key: diskspace|freespace
Diskspace|Capacity (GB)	Total capacity of datastore in gigabytes. Key: diskspace|capacity
Diskspace|Overhead	Amount of disk space that is overhead. Key: diskspace|overhead

Summary Metrics for Datastores

Summary metrics provide information about overall performance.

Metric	Description
* Summary|Total Number of Hosts	This metric shows the number of hosts that the datastore is connected to. Use this metric to know how many clusters the datastore is attached to. The number should not be too high, as a datastore should not be mounted by every host. The datastore and cluster should be paired to keep operations simple. Key: summary|total_number_hosts
* Summary|Total Number of VMs	This metric shows the number of virtual machines which save their VMDK files on the datastore. If a VM has four VMDKs stored in four datastores, the VM will be counted on each datastore. Use this metric to know how many VMs have at least one VMDK on a specific datastore. The number of VMs should be within your Concentration Risk policy. You should also expect the datastore to be well used. If only a few VMs are using the datastore, this is not considered a good use. Key: summary|total_number_vms
Summary|Maximum Number of VMs	Maximum number of virtual machines. Key: summary|max_number_vms
Summary|Workload Indicator	Workload indicator. Key: summary|workload_indicator
* Summary|Total Number of Clusters	This metric shows the number of clusters that the datastore is connected to. Key: summary|total_number_clusters

Template Metrics for Datastores

Metric	Description
Template|Virtual Machine used	Space used by virtual machine files. Key: template|used
Template|Access Time	Last access time. Key: template|accessTime

Calculated Metrics

vRealize Operations Manager calculates metrics for capacity, badges, and the health of the system. Calculated metrics apply to a subset of objects found in the `describe.xml` file that describes each adapter.

From data that the vCenter adapter collects, vRealize Operations Manager calculates metrics for objects of type:

- vSphere World

- Virtual Machine
- Host System
- Datastore

From data that the vRealize Operations Manager adapter collects, vRealize Operations Manager calculates metrics for objects of type:

- Node
- Cluster

Capacity and Project-Based Metrics

The capacity engine computes and publishes metrics that help you to plan your resource use based on consumer demand. Project-based metrics are a subset of capacity metrics that help to plan future resource use based on predicted consumer demand.

Capacity Metrics Group

For the capacity metrics group, full metric names include the name of the resource container. For example, if density metrics are computed for CPU or memory, the actual metric names appear as cpu|density or mem|density.

Only resource containers enabled for the capacity computations have relevant metrics. Not all metric types are generated for all resource containers. For example, if CPU or memory resource containers are enabled in a policy for density, but the network resource container is not, then cpu|density and mem|density metrics are calculated but network|density metrics are not.

A capacity metric definition includes resource containers that act as a consumer or a provider. For example in vSphere, the virtual machines are consumers of CPU and memory that the ESX host provides.

Table 7-38. Capacity Metrics Group

Metric Key	Metric Name	Generated for	Description
capacityRemainingUsingConsumers_average	Capacity Remaining for Average Consumer Profile	Provider	Number of average-size consumers that can fit into the capacity remaining. An average-size consumer demands 50% of total capacity.
capacityRemainingUsingConsumers_small	Capacity Remaining for Small Consumer Profile	Provider	Number of small-size consumers that can fit into the capacity remaining. A small-size consumer demands 0 - 33% of the total capacity.
capacityRemainingUsingConsumers_medium	Capacity Remaining for Medium Consumer Profile	Provider	Number of medium-size consumers that can fit into the capacity remaining. A medium-size consumer demands 33-66% of the total capacity.

Table 7-38. Capacity Metrics Group (continued)

Metric Key	Metric Name	Generated for	Description
capacityRemainingUsingConsumers_large	Capacity Remaining for Large Consumer Profile	Provider	Number of large-size consumers that can fit into the capacity remaining. A large-size consumer demands 66-100% of the total capacity.
capacityRemaining	Capacity Remaining (%)	Both	Percent capacity remaining in the resource container. For example, if the resource container is memory and 2 GB out of 10 GB of memory is free, the capacityRemaining = 20%.
underusedpercent	Under used (%)	Both	Percent capacity not being used.
idletimepercent	Idle time (%)	Both	Percent time a resource is idle based on use over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a resource is idle for a total of 6 days out of 30 days, idletimepercent = 20%.
wasteValue	Reclaimable Capacity	Both	Amount of reclaimable capacity based on consumer demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a vSphere host is configured with 10 GB of memory but only 2 GB of memory is used on average over 30 days, then wasteValue = 8 GB.
size.recommendation	Recommended Size	Both	Capacity recommendation based on demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if consumer demand is 2 GB of memory on average over 30 days, then the capacity recommendation is 2 GB.
optimal.vConsumption.per.pConsumption	Optimal consumption ratio	Provider	Ratio of ideal resource consumption to provision based on consumer demand over time. Ideal resource consumption is when the current capacity satisfies demand. Time is a policy setting. If not set, the default period is 30 days.
vConsumption.per.pConsumption	Consumption ratio	Provider	Ratio of current resource consumption to provision based on consumer demand.
object.demand	Stress Free Demand	Both	Demand based on peak analysis of raw demand values.
object.capacity	Usable Capacity	Both	Total capacity minus buffers. Capacity buffer is a policy setting.

Table 7-38. Capacity Metrics Group (continued)

Metric Key	Metric Name	Generated for	Description
object.demand.percent	Effective Demand (%)	Both	Percent capacity required by effective demand.
powered.on.consumer.count	Number of powered on consumers	Both	Number of consumers that are using a resource.
base.demand	Computed Demand	Both	Demand for an object based on self or consumer demand without the peak consideration policy setting.
actual.capacity	Current size	Both	Actual capacity without buffers
wastePercent	Reclaimable Capacity (%)	Both	Percent of reclaimable capacity based on consumer demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a vSphere host is configured with 10 GB of memory but only 2 GB of memory is used on average over 30 days, then wastePercent = 80%.

Object-level Metrics Group

Object-level metrics are calculated to track capacity use for all objects of a particular object type.

Table 7-39. Object-level Metrics Group

Metric Key	Metric Name	Description
summary|timeRemaining	Time Remaining	Time remaining before usable capacity runs out. Usable capacity excludes capacity reserved for HA and buffers.
summary|isStress	Is Stressed	Value equals 1 or a yellow badge indicates that an object is stressed. Value equals 0 or a green badge indicates that the object is not stressed. For a stress badge defined in a policy, when the stress exceeds the lowest threshold, the badge color changes from green to yellow.
summary|capacityRemainingValue	Capacity Remaining Value	Capacity remaining.
summary|oversized	Is Oversized	Indicates if an object has too much capacity configured, value of 1, or not, value of 0.
summary|idle	Is Idle	Indicates if an object is idle (value of 1) or not (value of 0).
summary|poweredOff	Powered Off	Indicates power state of an object. Value of 1 means ON and value of 0 means OFF.

Table 7-39. Object-level Metrics Group (continued)

Metric Key	Metric Name	Description
summary|capacityRemainingUsingConsumers_average	Capacity Remaining (Average consumer profile)	Capacity remaining based on average consumer demand.
summary|capacityRemainingUsingConsumers_small	Capacity Remaining (Small consumer profile)	Capacity remaining based on small consumer demand.
summary|capacityRemainingUsingConsumers_medium	Capacity Remaining (Medium consumer profile)	Capacity remaining based on medium consumer demand.
summary|capacityRemainingUsingConsumers_large	Capacity Remaining (Large consumer profile)	Capacity remaining based on large consumer demand.
summary|capacityRemaining_min	Capacity Remaining (Based on instantaneous peak)	Capacity remaining based on peak demand or stress.
summary|capacity.provider.count	Number of Capacity providers	Number of capacity providers.
summary|consumer.count	Number of Capacity consumers	Number of capacity consumers.
summary|consumer.count.per.provider.count	Consumer Provider ratio	Ratio of number of consumers to number of providers.
summary|optimal.consumer.per.provider	Optimal Consumer Provider ratio	Ratio of consumer to provider that would be optimal based on consumer demand.

Project-Based Metrics

Project-based metrics are calculated for a change in resources or demand that could affect capacity at some time in the future. See Chapter 6 Planning the Capacity for Your Managed Environment Using vRealize Operations Manager. Most metrics appear with `_whatif` appended to the capacity metric name. For example, the what-if applicable metric for capacity remaining is published as `capacityRemaining_whatif`.

Badge Metrics

Badge metrics provide information for badges in the user interface. They report the health, risk, and efficiency of objects in your environment.

vRealize Operations Manager 6.x analyzes badge metric data at five-minute averages, instead of hourly. As a result, you might find that efficiency and risk badge calculations are more sensitive than in previous versions. Badge metrics continue to be published nightly.

Table 7-40. Badge Metrics

Metric Key	Metric Name	Description
badge|alert_count_critical	Alert Count Critical	Count of critical alerts on the object.
badge|alert_count_immediate	Alert Count Immediate	Count of immediate alerts on the object.
badge|alert_count_info	Alert Count Info	Count of info alerts on the object.
badge| alert_count_warning	Alert Count Warning	Count of warning alerts on the object.
badge|anomaly	Anomaly	Overall score for anomalies, on a scale of 100.
badge|capacityRemaining	Capacity Remaining	Overall score for capacity remaining, on a scale of 100.
badge|compliance	Compliance	Overall score for compliance, on a scale of 100.
badge|density	Density	Overall score for density, on a scale of 100.
badge|efficiency	Efficiency	Overall score for efficiency. The score will be one of these discrete values representing each state of the badge: Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1.
badge|efficiency_classic	Legacy Efficiency	The legacy efficiency score computed on a scale of 100 as per vCenter Operations Manager version 5.x. For backward compatibility purposes.
badge|efficiency_state	Efficiency State	Represents the state of the efficiency badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1.
badge|fault	Fault	Overall score for fault, on a scale of 100.
badge|health	Health	Overall score for health. The score will be one of these discrete values representing each state of the badge: Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1.
badge|health_classic	Legacy Health	The legacy health score computed on a scale of 100 as per vCenter Operations Manager 5.x. For backward compatibility purposes.
badge|health_state	Health State	Represents the state of health badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1
badge|risk	Risk	Overall score for risk. The score will be one of these discrete values representing each state of the badge: Green - 0, Yellow - 25, Orange - 50, Red - 75, Unknown: -1.
badge|risk_classic	Legacy Risk	The legacy risk score computed on a scale of 100 as per vCenter Operations Manager 5.x. For backward compatibility purposes.
badge|risk_state	Risk State	Represents the state of risk badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1.
badge|stress	Stress	Overall score of stress, on a scale of 100.

Table 7-40. Badge Metrics (continued)

Metric Key	Metric Name	Description
badge|timeRemaining	Time Remaining - Real Time	Overall score of real time remaining, on a scale of 100.
badge|waste	Waste	Overall score of waste, on a scale of 100.
badge|workload	Workload (%)	Overall score of workload, on a scale of 100.

System Metrics

System metrics provide information used to monitor the health of the system. They help you to identify problems in your environment.

Table 7-41. System Metrics

Metric Key	Metric Name	Description
System Attributes|health	Self - Health Score	System health score of self resource
System Attributes|all_metrics	Self - Metric Count	Number of metrics of self resource
System Attributes|ki_metrics	Self - KPI Count	Number of KPI metrics of self resource
System Attributes|active_alarms	Self - Active Anomaly Count	Number of active alarms of self resource
System Attributes|new_alarms	Self - New Anomaly Count	Number of new alarms of self resource
System Attributes|active_ki_alarms	Self - Active KPI Breach Count	Number of active KPI alarms of self resource
System Attributes|new_ki_alarms	Self - New KPI Breach Count	Number of new KPI alarms of self resource
System Attributes|total_alarms	Self - Total Anomalies	Number of total alarms of self resource
System Attributes|change_index	Self - Change Index	Change index of self resource(100 - health score)
System Attributes|child_all_metrics	Full Set - Metric Count	Number of metrics of child resources
System Attributes|child_ki_metrics	Full Set - KPI Count	Number of KPI metrics of child resources
System Attributes|child_active_alarms	Full Set - Active Anomaly Count	Number of active alarms of child resources
System Attributes|child_new_alarms	Full Set - New Anomaly Count	Number of new alarms of child resources
System Attributes| child_active_ki_alarms	Full Set - Active KPI Breach Count	Number of active KPI alarms of child resources
System Attributes|child_new_ki_alarms	Full Set - New KPI Breach Count	Number of new KPI alarms of child resources
System Attributes|availability	Availability	Resource availability (0-down, 1-Up, -1-Unknown)
System Attributes|alert_count_critical	Alert Count Critical	Number of Critical alerts

Table 7-41. System Metrics (continued)

Metric Key	Metric Name	Description
System Attributes| alert_count_immediate	Alert Count Immediate	Number of Immediate alerts
System Attributes|alert_count_warning	Alert Count Warning	Number of Warning alerts
System Attributes|alert_count_info	Alert Count Info	Number of Info alerts

Self-Monitoring Metrics for vRealize Operations Manager

vRealize Operations Manager uses the vRealize Operations Manager adapter to collect metrics that monitor its own performance. These self-monitoring metrics drive capacity models for vRealize Operations Manager objects and are useful for diagnosing problems with vRealize Operations Manager.

Analytics Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager analytics service, including threshold checking metrics.

Table 7-42. Analytics Metrics

Metric Key	Metric Name	Description
ActiveAlarms	Active DT Symptoms	Active DT Symptoms.
ActiveAlerts	Active Alerts	Active alerts.
PrimaryResourcesCount	Number of primary objects	Number of primary objects
LocalResourcesCount	Number of local objects	Number of local objects
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
LocalMetricsCount	Number of local metrics	Number of local metrics
ReceivedResourceCount	Number of received objects	Number of received objects
ReceivedMetricCount	Number of received metrics	Number of received metrics
LocalFDSize	Number of forward data entries	Number of locally stored primary and redundant entries in forward data region.
LocalPrimaryFDSize	Number of primary forward data entries	Number of locally stored primary entries in forward data region.
LocalFDAltSize	Number of alternative forward data entries	Number of locally stored primary and redundant entries in alternative forward data region.
LocalPrimaryFDAltSize	Number of alternative primary forward data entries	Number of locally stored primary entries in alternative forward data region.
CurrentHeapSize	Current heap size	Current heap size.
MaxHeapSize	Max heap size	Max heap size
CommittedMemory	Committed memory	Committed memory

Table 7-42. Analytics Metrics (continued)

Metric Key	Metric Name	Description
CPUUsage	CPU usage	CPU usage
Threads	Threads	Threads
UpStatus	Threads	Threads

Overall Threshold Checking Metrics for the Analytics Service

Overall threshold checking captures various metrics for work items used to process incoming observation data. All metrics keys for the overall threshold checking metrics begin with OverallThresholdChecking, as in `OverallThresholdChecking|Count` or `OverallThresholdChecking|CheckThresholdAndHealth|OutcomeObservationsSize|TotalCount`.

Table 7-43. Overall Threshold Checking Metrics for the Analytics Service

Metric Key	Metric Name	Description
Count	Count	Count
Duration|TotalDuration	Total	Total length of duration (ms)
Duration|AvgDuration	Average	Average duration (ms)
Duration|MinDuration	Minimum	Minimum duration (ms)
Duration|MaxDuration	Maximum	Maximum duration (ms)
IncomingObservationsSize|TotalCount	Total	Total
IncomingObservationsSize|AvgCount	Average	Average
IncomingObservationsSize|MinCount	Minimal	Minimal
IncomingObservationsSize|MaxCount	Maximal	Maximal
CheckThresholdAndHealth|Count	Count	Count
CheckThresholdAndHealth|Duration|TotalDuration	Total	Total length of duration (ms)
CheckThresholdAndHealth|Duration|AvgDuration	Average	Average duration (ms)
CheckThresholdAndHealth|Duration|MinDuration	Minimum	Minimum duration (ms)
CheckThresholdAndHealth|Duration|MaxDuration	Maximum	Maximum duration (ms)
CheckThresholdAndHealth|OutcomeObservationsSize|TotalCount	Total	Total
CheckThresholdAndHealth|OutcomeObservationsSize|AvgCount	Average	Average
CheckThresholdAndHealth|OutcomeObservationsSize|MinCount	Minimal	Minimal
CheckThresholdAndHealth|OutcomeObservationsSize|MaxCount	Maximal	Maximal

Table 7-43. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
SuperMetricComputation|Count	Count	Count
SuperMetricComputation|Duration|TotalDuration	Total	Total length of duration (ms)
SuperMetricComputation|Duration|AvgDuration	Average	Average duration (ms)
SuperMetricComputation|Duration|MinDuration	Minimum	Minimum duration (ms)
SuperMetricComputation|Duration|MaxDuration	Maximum	Maximum duration (ms)
SuperMetricComputation|SuperMetricsCount|TotalCount	Total	Total
SuperMetricComputation|SuperMetricsCount|AvgCount	Average	Average
SuperMetricComputation|SuperMetricsCount|MinCount	Minimal	Minimal
SuperMetricComputation|SuperMetricsCount|MaxCount	Maximal	Maximal
StoreObservationToFSDB|Count	Count	Count
StoreObservationToFSDB|Duration|TotalDuration	Total	Total length of duration (ms)
StoreObservationToFSDB|Duration|AvgDuration	Average	Average duration (ms)
StoreObservationToFSDB|Duration|MinDuration	Minimum	Minimum duration (ms)
StoreObservationToFSDB|Duration|MaxDuration	Maximum	Maximum duration (ms)
StoreObservationToFSDB|StoredObservationsSize|TotalCount	Total	Total
StoreObservationToFSDB|StoredObservationsSize|AvgCount	Average	Average
StoreObservationToFSDB|StoredObservationsSize|MinCount	Minimal	Minimal
StoreObservationToFSDB|StoredObservationsSize|MaxCount	Maximal	Maximal
UpdateResourceCache|Count	Count	Count
UpdateResourceCache|Duration|TotalDuration	Total	Total
UpdateResourceCache|Duration|AvgDuration	Average	Average
UpdateResourceCache|Duration|MinDuration	Minimum	Minimum

Table 7-43. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
UpdateResourceCache|Duration|MaxDuration	Maximum	Maximum
UpdateResourceCache|ModifcationEstimateCount|TotalCount	Total	The number of estimated modifications done during each resource cache object update.
UpdateResourceCache|ModifcationEstimateCount|AvgCount	Average	Average
UpdateResourceCache|ModifcationEstimateCount|MinCount	Minimal	Minimal
UpdateResourceCache|ModifcationEstimateCount|MaxCount	Maximal	Maximal
ManageAlerts|Count	Count	The total number of times the threshold checking work items perform alert updates.
ManageAlerts|Duration|TotalDuration	Total	The duration for the alert updates operations.
ManageAlerts|Duration|AvgDuration	Average	Average
ManageAlerts|Duration|MinDuration	Minimum	Minimum
ManageAlerts|Duration|MaxDuration	Maximum	Maximum
UpdateSymptoms|Count	Count	The total number of times the threshold checking work items check and build symptoms.
UpdateSymptoms|Duration|TotalDuration	Total	The duration for the check and build symptoms operation.
UpdateSymptoms|Duration|AvgDuration	Average	Average
UpdateSymptoms|Duration|MinDuration	Minimum	Minimum
UpdateSymptoms|Duration|MaxDuration	Maximum	Maximum

Dynamic Threshold Calculation Metrics for the Analytics Service

All metrics keys for the dynamic threshold calculation metrics begin with DtCalculation, as in `DtCalculation|DtDataWrite|WriteOperationCount` or `DtCalculation|DtAnalyze|AnalyzeOperationCount`.

Table 7-44. Dynamic Threshold Calculation Metrics for the Analytics Service

Metric Key	Metric Name	Description
DtDataWrite|WriteOperationCount	Write operation count	Write operation count
DtDataWrite|Duration|TotalDuration	Total	Total length of duration (ms)
DtDataWrite|Duration|AvgDuration	Average	Average duration (ms)
DtDataWrite|Duration|MinDuration	Minimum	Minimum duration (ms)
DtDataWrite|Duration|MaxDuration	Maximum	Maximum duration (ms)

Table 7-44. Dynamic Threshold Calculation Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
DtDataWrite|SavedDtObjectCount|TotalCount	Total	Total
DtDataWrite|SavedDtObjectCount|AvgCount	Average	Average
DtDataWrite|SavedDtObjectCount|MinCount	Minimal	Minimal
DtDataWrite|SavedDtObjectCount|MaxCount	Maximal	Maximal
DtAnalyze|AnalyzeOperationCount	Analyze Operation Count	Analyze Operation Count
DtAnalyze|Duration|TotalDuration	Total	Total length of duration (ms)
DtAnalyze|Duration|AvgDuration	Average	Average duration (ms)
DtAnalyze|Duration|MinDuration	Minimum	Minimum duration (ms)
DtAnalyze|Duration|MaxDuration	Maximum	Maximum duration (ms)
DtAnalyze|AnalyzedMetricsCount|TotalCount	Total	Total
DtAnalyze|AnalyzedMetricsCount|AvgCount	Average	Average
DtAnalyze|AnalyzedMetricsCount|MinCount	Minimal	Minimal
DtAnalyze|AnalyzedMetricsCount|MaxCount	Maximal	Maximal
DtDataRead|ReadOperationsCount	Read Operation Count	Read Operation Count
DtDataRead|Duration|TotalDuration	Total	Total length of duration (ms)
DtDataRead|Duration|AvgDuration	Average	Average duration (ms)
DtDataRead|Duration|MinDuration	Minimum	Minimum duration (ms)
DtDataRead|Duration|MaxDuration	Maximum	Maximum duration (ms)
DtDataRead|ReadDataPointsCount|TotalCount	Total	Total
DtDataRead|ReadDataPointsCount|AvgCount	Average	Average
DtDataRead|ReadDataPointsCount|MinCount	Minimal	Minimal
DtDataRead|ReadDataPointsCount|MaxCount	Maximal	Maximal

Table 7-45. Function Call Metrics for the Analytics Service

Metric Key	Metric Name	Description
FunctionCalls|Count	Number of function calls	Number of function calls
FunctionCalls|AvgDuration	Average execution time	Average execution time
FunctionCalls|MaxDuration	Max execution time	Max execution time

Collector Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager Collector service objects.

Table 7-46. Collector Metrics

Metric Key	Metric Name	Description
ThreadpoolThreadsCount	Number of pool threads	Number of pool threads.
RejectedFDCount	Number of rejected forward data	Number of rejected forward data
RejectedFDAltCount	Number of rejected alternative forward data	Number of rejected alternative forward data
SentFDCount	Number of sent objects	Number of sent objects
SentFDAltCount	Number of alternative sent objects	Number of alternative sent objects
CurrentHeapSize	Current heap size (MB)	Current heap size.
MaxHeapsize	Max heap size (MB)	Maximum heap size.
CommittedMemory	Committed memory (MB)	Amount of committed memory.
CPUUsage	CPU usage	CPU usage.
Threads	Threads	Number of threads.
UpStatus	Up Status	Up Status

Controller Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager Controller objects.

Table 7-47. Controller Metrics

Metric Key	Metric Name	Description
RequestedMetricCount	Number of requested metrics	Number of requested metrics
ApiCallsCount	Number of API calls	Number of API calls
NewDiscoveredResourcesCount	Number of discovered objects	Number of discovered objects

FSDB Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager file system database (FSDB) objects.

Table 7-48. FSDB Metrics

Metric Key	Metric Name	Description
StoragePoolElementsCount	Number of storage work items	Number of storage work items
FsdbState	Fsdb state	Fsdb state
StoredResourcesCount	Number of stored objects	Number of stored objects
StoredMetricsCount	Number of stored metrics	Number of stored metrics

Table 7-49. Storage Thread Pool Metrics for FSDB

Metric Key	Metric Name	Description
StoreOperationsCount	Store operations count	Store operations count
StorageThreadPool|Duration| TotalDuration	Total	Total number of duration (ms)
StorageThreadPool|Duration|AvgDuration	Average	Average duration (ms)
StorageThreadPool|Duration|MinDuration	Minimum	Minimum duration (ms)
StorageThreadPool|Duration|MaxDuration	Maximum	Maximum duration (ms)
StorageThreadPool|SavedMetricsCount| TotalCount	Total	Total
StorageThreadPool|SavedMetricsCount| AvgCount	Average	Average
StorageThreadPool|SavedMetricsCount| MinCount	Minimal	Minimal
StorageThreadPool|SavedMetricsCount| MaxCount	Maximal	Maximal

Product UI Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager product user interface objects.

Table 7-50. Product UI Metrics

Metric Key	Metric Name	Description
ActiveSessionsCount	Active sessions	Active sessions
CurrentHeapSize	Current heap size	Current heap size.
MaxHeapsize	Max heap size	Maximum heap size.
CommittedMemory	Committed memory	Amount of committed memory.
CPUUsage	CPU usage	Percent CPU use.
Threads	Threads	Number of threads.
SessionCount	Number of active sessions	Number of active sessions
SelfMonitoringQueueSize	Self Monitoring queue size	Self Monitoring queue size

Table 7-51. API Call Metrics for the Product UI

Metric Key	Metric Name	Description
APICalls|HTTPRequesterRequestCount	HTTPRequester request count	HTTPRequester request count
APICalls|AvgHTTPRequesterRequestTime	HTTPRequester average request time	HTTPRequester average request time (ms)
APICalls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
APICalls|AvgAlertRequestTime	Average alert request time	Average alert request time (ms)
APICalls|AlertRequestCount	Alert request count	Alert request count

Table 7-51. API Call Metrics for the Product UI (continued)

Metric Key	Metric Name	Description
APICalls|AvgMetricPickerRequestTime	Average metric-picker request time	Average metric-picker request time (ms)
APICalls|MetricPickerRequestCount	Metric picker request count	Metric picker request count
APICalls|HeatmapRequestCount	Heatmap request count	Heatmap request count
APICalls|AvgHeatmapRequestTime	Average HeatMap request time	Average HeatMap request time (ms)
APICalls|MashupChartRequestCount	Mashup Chart request count	Mashup Chart request count
APICalls|AvgMashupChartRequestTime	Average Mashup Chart request time	Average Mashup Chart request time (ms)
APICalls|TopNRequestCount	Top N request count	Top N request count
APICalls|AvgTopNRequestTime	Average Top N request time	Average Top N request time (ms)
APICalls|MetricChartRequestCount	Metric Chart request count	Metric Chart request count
APICalls|AvgMetricChartRequestTime	Average MetricChart request time	Average MetricChart request time (ms)

Admin UI Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager administration user interface objects.

Table 7-52. Admin UI Metrics

Metric Key	Metric Name	Description
CurrentHeapSize	Current heap size	Current heap size (MB).
MaxHeapsize	Max heap size	Maximum heap size (MB).
CommittedMemory	Committed memory	Amount of committed memory (MB) .
CPUUsage	CPU usage	CPU usage (%).
Threads	Threads	Number of threads.
SessionCount	Number of active sessions	Number of active sessions
SelfMonitoringQueueSize	Self Monitoring queue size	Self Monitoring queue size

Table 7-53. API Call Metrics for the Admin UI

Metric Key	Metric Name	Description
APICalls|HTTPRequesterRequestCount	HTTPRequester request count	HTTPRequester request count
APICalls|AvgHTTPRequesterRequestTime	HTTPRequester average request time	HTTPRequester average request time (ms)

Suite API Metrics

vRealize Operations Manager collects metrics for the VMware vRealize Operations Management Suite API objects.

Table 7-54. Suite API Metrics

Metric Key	Metric Name	Description
UsersCount	Number of users	Number of users
ActiveSessionsCount	Active sessions	Active sessions
GemfireClientReconnects	Gemfire Client Reconnects	Gemfire Client Reconnects
GemfireClientCurrentCalls	Gemfire Client Total Outstanding	Gemfire Client Total Outstanding
CurrentHeapSize	Current heap size	Current heap size (MB) .
MaxHeapsize	Max heap size	Maximum heap size (MB) .
CommittedMemory	Committed memory	Amount of committed memory (MB).
CPUUsage	CPU usage	CPU usage (%) .
CPUProcessTime	CPU process time	CPU process time (ms)
CPUProcessTimeCapacity	CPU process time capacity	CPU process time capacity (ms)
Threads	Threads	Number of threads.

Table 7-55. Gemfire Client Call Metrics for the Suite API

Metric Key	Metric Name	Description
GemfireClientCalls|TotalRequests	Total Requests	Total Requests
GemfireClientCalls|AvgResponseTime	Average Response Time	Average Response Time (ms)
GemfireClientCalls|MinResponseTime	Minimum Response Time	Minimum Response Time (ms)
GemfireClientCalls|MaxResponseTime	Maximum Response Time	Maximum Response Time
GemfireClientCalls|RequestsPerSecond	Requests per Second	Requests per Second
GemfireClientCalls|CurrentRequests	Current Requests	Current Requests
GemfireClientCalls|RequestsCount	Requests Count	Requests Count
GemfireClientCalls|ResponsesCount	Responses Count	Responses Count

Table 7-56. API Call Metrics for the Suite API

Metric Key	Metric Name	Description
APICalls|TotalRequests	Total Requests	Total Requests
APICalls|AvgResponseTime	Average Response Time (ms)	Average Response Time (ms)
APICalls|MinResponseTime	Minimum Response Time (ms)	Minimum Response Time (ms)
APICalls|MaxResponseTime	Maximum Response Time	Maximum Response Time
APICalls|ServerErrorResponseCount	Server Error Response Count	Server Error Response Count
APICalls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
APICalls|FailedAuthorizationCount	Failed Authorization Count	Failed Authorization Count
APICalls|RequestsPerSecond	Requests per Second	Requests per Second
APICalls|CurrentRequests	Current Requests	Current Requests
APICalls|ResponsesPerSecond	Responses per Second	Responses per Second

Table 7-56. API Call Metrics for the Suite API (continued)

Metric Key	Metric Name	Description
APICalls|RequestsCount	Requests Count	Requests Count
APICalls|ResponsesCount	Responses Count	Responses Count

Cluster and Slice Administration Metrics

vRealize Operations Manager collects metrics for vRealize Operations Manager Cluster and Slice Administration (CaSA) objects.

Table 7-57. Cluster and Slice Administration Metrics

Metric Key	Metric Name	Description
CurrentHeapSize	Current heap size	Current heap size (MB).
MaxHeapsize	Max heap size	Maximum heap size (MB).
CommittedMemory	Committed memory	Amount of committed memory (MB).
CPUUsage	CPU usage	CPU usage (%)
Threads	Threads	Number of threads.

Table 7-58. API Call Metrics for Cluster and Slice Administration

Metric Key	Metric Name	Description
API Calls|TotalRequests	Total Requests	Total Requests
API Calls|AvgResponseTime	Average Response Time	Average Response Time (ms)
API Calls|MinResponseTime	Minimum Response Time	Minimum Response Time (ms)
API Calls|MaxResponseTime	Maximum Response Time	Maximum Response Time (ms)
API Calls|ServerErrorResponseCount	Server Error Response Count	Server Error Response Count
API Calls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
API Calls|FailedAuthorizationCount	Minimum Response Time	Minimum Response Time (ms)

Watchdog Metrics

vRealize Operations Manager collects watchdog metrics to ensure that the vRealize Operations Manager services are running and responsive.

Watchdog Metrics

The watchdog metric provides the total service count.

Table 7-59. Watchdog Metrics

Metric Key	Metric Name	Description
ServiceCount	Service Count	Service Count

Service Metrics

Service metrics provide information about watchdog activity.

Table 7-60. Metrics for the vRealize Operations Manager Watchdog Service

Metric Key	Metric Name	Description
Service|Enabled	Enabled	Enabled
Service|Restarts	Restarts	Number of times the process has been unresponsive and been restarted by Watchdog.
Service|Starts	Starts	Number of times the process has been revived by Watchdog.
Service|Stops	Stops	Number of times the process has been stopped by Watchdog.

Node Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager node objects.

Metrics can be calculated for node objects. See Calculated Metrics.

Table 7-61. Node Metrics

Metric Key	Metric Name	Description
Component Count	Component count	The number of vRealize Operations Manager objects reporting for this node
PrimaryResourcesCount	Number of primary objects	Number of primary objects
LocalResourcesCount	Number of local objects	Number of local objects
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
LocalMetricsCount	Number of local metrics	Number of local metrics
PercentDBStorageAvailable	Percent disk available /storage/db	Percent disk available /storage/db
PercentLogStorageAvailable	Percent disk available /storage/log	Percent disk available /storage/log

Table 7-62. Memory Metrics for the Node

Metric Key	Metric Name	Description
mem|actualFree	Actual Free	Actual Free
mem|actualUsed	Actual Used	Actual Used
mem|free	Free	Free)
mem|used	Used	Used
mem|total	Total	Total
mem|demand_gb	Estimated memory demand	Estimated memory demand

Table 7-63. Swap Metrics for the Node

Metric Key	Metric Name	Description
swap|total	Total	Total
swap|free	Free	Free
swap|used	Used	Used
swap|pageIn	Page in	Page in
swap|pageOut	Page out	Page out

Table 7-64. Resource Limit Metrics for the Node

Metric Key	Metric Name	Description
resourceLimit|numProcesses	Number of processes	Number of processes
resourceLimit|openFiles	Number of open files	Number of open files
resourceLimit|openFilesMax	Number of open files maximum limit	Number of open files maximum limit
resourceLimit|numProcessesMax	Number of processes maximum limit	Number of processes maximum limit

Table 7-65. Network Metrics for the Node

Metric Key	Metric Name	Description
net|allInboundTotal	All inbound connections	All inbound total
net|allOutboundTotal	All outbound connections	All outbound total
net|tcpBound	TCP bound	TCP bound
net|tcpClose	TCP state CLOSE	Number of connections in TCP CLOSE
net|tcpCloseWait	TCP state CLOSE WAIT	Number of connections in TCP state CLOSE WAIT
net|tcpClosing	TCP state CLOSING	Number of connections in TCP state CLOSING
net|tcpEstablished	TCP state ESTABLISHED	Number of connections in TCP state ESTABLISHED
net|tcpIdle	TCP state IDLE	Number of connections in TCP state IDLE
net|tcpInboundTotal	TCP inbound connections	TCP inbound connections
net|tcpOutboundTotal	TCP outbound connections	TCP outbound connections
net|tcpLastAck	TCP state LAST ACK	Number of connections in TCP state LAST ACK
net|tcpListen	TCP state LISTEN	Number of connections in TCP state LISTEN
net|tcpSynRecv	TCP state SYN RCVD	Number of connections in TCP state SYN RCVD

Table 7-65. Network Metrics for the Node (continued)

Metric Key	Metric Name	Description
net|tcpSynSent	TCP state SYN_SENT	Number of connections in TCP state SYN_SENT
net|tcpTimeWait	TCP state TIME WAIT	Number of connections in TCP state TIME WAIT

Table 7-66. Network Interface Metrics for the Node

Metric Key	Metric Name	Description
net|iface|speed	Speed	Speed (bits/sec)
net|iface|rxPackets	Receive packets	Number of received packets
net|iface|rxBytes	Receive bytes	Number of received bytes
net|iface|rxDropped	Receive packet drops	Number of received packets dropped
net|iface|rxFrame	Receive packets frame	Number of receive packets frame
net|iface|rxOverruns	Receive packets overruns	Number of receive packets overrun
net|iface|txPackets	Transmit packets	Number of transmit packets
net|iface|txBytes	Transmit bytes	Number of transmit bytes
net|iface|txDropped	Transmit packet drops	Number of transmit packets dropped
net|iface|txCarrier	Transmit carrier	Transmit carrier
net|iface|txCollisions	Transmit packet collisions	Number of transmit collisions
net|iface|txErrors	Transmit packet errors	Number of transmit errors
net|iface|txOverruns	Transmit packet overruns	Number of transmit overruns

Table 7-67. Disk Filesystem Metrics for the Node

Metric Key	Metric Name	Description
disk|fileSystem|total	Total	Total
disk|fileSystem|available	Available	Available
disk|fileSystem|used	Used	Used
disk|fileSystem|files	Total file nodes	Total file nodes
disk|fileSystem|filesFree	Total free file nodes	Total free file nodes
disk|fileSystem|queue	Disk queue	Disk queue
disk|fileSystem|readBytes	Read bytes	Number of bytes read
disk|fileSystem|writeBytes	Write bytes	Number of bytes written
disk|fileSystem|reads	Reads	Number of reads
disk|fileSystem|writes	Writes	Number of writes

Table 7-68. Disk Installation Metrics for the Node

Metric Key	Metric Name	Description
disk|installation|used	Used	Used
disk|installation|total	Total	Total
disk|installation|available	Available	Available

Table 7-69. Disk Database Metrics for the Node

Metric Key	Metric Name	Description
disk|db|used	Used	Used
disk|db|total	Total	Total
disk|db|available	Available	Available

Table 7-70. Disk Log Metrics for the Node

Metric Key	Metric Name	Description
disk|log|used	Used	Used
disk|log|total	Total	Total
disk|log|available	Available	Available

Table 7-71. CPU Metrics for the Node

Metric Key	Metric Name	Description
cpu|combined	Combined load	Combined load (User + Sys + Nice + Wait)
cpu|idle	Idle	Idle time fraction of total available cpu (cpu load)
cpu|irq	Irq	Interrupt time fraction of total available cpu (cpu load)
cpu|nice	Nice	Nice time fraction of total available cpu (cpu load)
cpu|softIrq	Soft Irq	Soft interrupt time fraction of total available cpu (cpu load)
cpu|stolen	Stolen	Stolen time fraction of total available cpu (cpu load)
cpu|sys	Sys	Sys time fraction of total available cpu (cpu load)
cpu|user	User (cpu load)	User time fraction of total available cpu (cpu load)
cpu|wait	Wait (cpu load)	Wait time fraction of total available cpu (cpu load)
cpu|total	Total available for a cpu	Total available for a cpu
cpu|allCpuCombined	Total combined load for all cpus	Total combined load for all cpus (cpu load)

Table 7-71. CPU Metrics for the Node (continued)

Metric Key	Metric Name	Description
cpu|allCpuTotal_ghz	Available	Available
cpu|allCpuCombined_ghz	Used	Used
cpu|allCpuCombined_percent	CPU usage	CPU usage (%)

Table 7-72. Device Metrics for the Node

Metric Key	Metric Name	Description
device|iops	Reads/Writes per second	Average number of read/write commands issued per second during the collection interval.
device|await	Average transaction time	Average transaction time (milliseconds).
device|iops_readMaxObserved	Maximum observed reads per second	Maximum observed reads per second.
device|iops_writeMaxObserved	Maximum observed writes per second	Maximum observed writes per second.

Table 7-73. Service Metrics for the Node

Metric Key	Metric Name	Description
service|proc|fdUsage	Total number of open file descriptors	Total number of open file descriptors.

Table 7-74. NTP Metrics for the Node

Metric Key	Metric Name	Description
ntp|serverCount	Configured server count	Configured server count
ntp|unreachableCount	Unreachable server count	Unreachable server count
ntp|unreachable	Unreachable	Is the NTP server unreachable. Value of 0 is reachable, 1 means the server was not reached or didn't respond.

Table 7-75. Heap Metrics for the Node

Metric Key	Metric Name	Description
heap|CurrentHeapSize	Current heap size	Current heap size
heap|MaxHeapSize	Max heap size	Max heap size
heap|CommittedMemory	Committed Memory	Committed Memory

Cluster Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager cluster objects including dynamic threshold calculation metrics and capacity computation metrics.

Metrics can be calculated for cluster objects. See Calculated Metrics.

Cluster Metrics

Cluster metrics provide host, resource, and metric counts on the cluster.

Table 7-76. Cluster Metrics

Metric Key	Metric Name	Description
HostCount	Number of Nodes in Cluster	Number of Nodes in Cluster
PrimaryResourcesCount	Number of primary resources	Number of primary resources
LocalResourcesCount	Number of local resources	Number of local resources
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
ReceivedResourceCount	Number of received resources	Number of received resources
ReceivedMetricCount	Number of received metrics	Number of received metrics

DT Metrics

DT metrics are dynamic threshold metrics for the cluster. Non-zero values appear only if metric collection occurs while the dynamic threshold calculations are running.

Table 7-77. DT Metrics for the Cluster

Metric Key	Metric Name	Description
dt|isRunning	Running	Running
dt|dtRunTime	Running duration	Running duration (ms)
dt|StartTime	Running start time	Running start time
dt|percentage	Percent	Percent (%)
dt|executorCount	Executor Node Count	Executor Node Count
dt|resourceCount	Resource Count	Resource Count
dt|fsdbReadTime	FSDB Read Time	FSDB Read Time (ms)
dt|dtObjectSaveTime	DT Object Save Time	DT Object Save Time (ms)
dt|dtHistorySaveTime	DT History Save Time	DT History Save Time (ms)
dt|executor|resourceCount	Resource Count	Resource Count

Capacity Computation (CC) Metrics

CC metrics are capacity computation metrics for the cluster. Non-zero values appear only if metric collection occurs while the capacity computation calculations are running.

Table 7-78. CC Metrics for the Cluster

Metric Key	Metric Name	Description
cc|isRunning	Running	Running
cc|runTime	Total Run Time	Total Run Time
cc|startTime	Start time	Start time

Table 7-78. CC Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
cc|finishTime	Finish Time	Finish Time
cc|totalResourcesToProcess	Total Objects Count	Total Objects Count
cc|progress	Progress	Progress
cc|phase1TimeTaken	Phase 1 Computation Time	Phase 1 Computation Time
cc|phase2TimeTaken	Phase 2 Computation Time	Phase 2 Computation Time

Gemfire Cluster Metrics

Gemfire metrics provide information about the Gemfire cluster.

Table 7-79. Gemfire cluster Metrics for the Cluster

Metric Key	Metric Name	Description
GemfireCluster|System|AvgReads	Average reads per second	The average number of reads per second for all members
GemfireCluster|System|AvgWrites	Average writes per second	The average number of writes per second for all members
GemfireCluster|System|DiskReadsRate	Disk reads rate	The average number of disk reads per second across all distributed members
GemfireCluster|System|DiskWritesRate	Disk writes rate	The average number of disk writes per second across all distributed members
GemfireCluster|System| GarbageCollectionCount	Total garbage collection count	The total garbage collection count for all members
GemfireCluster|System| GarbageCollectionCountDelta	New garbage collection count	The new garbage collection count for all members
GemfireCluster|System|JVMPauses	JVM pause count	The number of detected JVM pauses
GemfireCluster|System|JVMPausesDelta	New JVM pause count	The number of new detected JVM pauses
GemfireCluster|System| DiskFlushAvgLatency	Disk flush average latency	Disk flush average latency (msec)
GemfireCluster|System| NumRunningFunctions	Number of running functions	The number of map-reduce jobs currently running on all members in the distributed system
GemfireCluster|System|NumClients	Number of clients	The number of connected clients
GemfireCluster|System|TotalHitCount	Total hit count	Total number of cache hits for all regions
GemfireCluster|System|TotalHitCountDelta	New hit count	Number of new cache hits for all regions
GemfireCluster|System|TotalMissCount	Total miss count	The total number of cache misses for all regions

Table 7-79. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
GemfireCluster|System|TotalMissCountDelta	New miss count	Number of new cache misses for all regions
GemfireCluster|System|Member|FreeSwapSpace	Swap space free	Swap space free (MB)
GemfireCluster|System|Member|TotalSwapSpace	Swap space total	Swap space total (MB)
GemfireCluster|System|Member|CommittedVirtualMemorySize	Committed virtual memory size	Committed virtual memory size (MB)
GemfireCluster|System|Member|SystemLoadAverage	System load average	System load average
GemfireCluster|System|Member|FreePhysicalMemory	Free physical memory	Free physical memory (MB)
GemfireCluster|System|Member|TotalPhysicalMemory	Total physical memory	Total physical memory (MB)
GemfireCluster|System|Member|CacheListenerCallsAvgLatency	Average cache listener calls latency	Average cache listener calls latency (msec)
GemfireCluster|System|Member|CacheWriterCallsAvgLatency	Average cache writer calls latency	Average cache writer calls latency (msec)
GemfireCluster|System|Member|DeserializationAvgLatency	Average deserialization latency	Average deserialization latency (msec)
GemfireCluster|System|Member|FunctionExecutionRate	Function executions per second	Function executions per second
GemfireCluster|System|Member|JVMPauses	Number of JVM pauses	Number of JVM pauses
GemfireCluster|System|Member|NumRunningFunctions	Number of running functions	Number of running functions
GemfireCluster|System|Member|PutsRate	Puts per second	Puts per second
GemfireCluster|System|Member|GetsRate	Gets per second	Gets per second
GemfireCluster|System|Member|GetsAvgLatency	Average gets latency	Average gets latency (msec)
GemfireCluster|System|Member|PutsAvgLatency	Average puts latency	Average puts latency (msec)
GemfireCluster|System|Member|SerializationAvgLatency	Average serialization latency	Average serialization latency (msec)
GemfireCluster|System|Member|Disk|DiskFlushAvgLatency	Flush average latency	Flush average latency (msec)
GemfireCluster|System|Member|Disk|DiskReadsRate	Average reads per second	Average reads per second
GemfireCluster|System|Member|Disk|DiskWritesRate	Average writes per second	Average writes per second

Table 7-79. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
GemfireCluster|System|Member|Network|BytesReceivedRate	Average received bytes per second	Average received bytes per second
GemfireCluster|System|Member|Network|BytesSentRate	Average sent bytes per second	Average sent bytes per second
GemfireCluster|System|Member|JVM|GCTimeMillis	Garbage Collection time	Total amount of time spent on garbage collection
GemfireCluster|System|Member|JVM|GCTimeMillisDelta	New Garbage Collection time	New amount of time spent on garbage collection
GemfireCluster|System|Member|JVM|TotalThreads	Total threads	Total threads
GemfireCluster|System|Member|JVM|CommitedMemory	Committed Memory	Committed Memory (MB)
GemfireCluster|System|Member|JVM|MaxMemory	Max Memory	Max Memory (MB)
GemfireCluster|System|Member|JVM|UsedMemory	Used Memory	Used Memory (MB)
GemfireCluster|Region|SystemRegionEntryCount	Entry Count	Entry Count
GemfireCluster|Region|DestroyRate	Destroys per second	Destroys per second
GemfireCluster|Region|CreatesRate	Creates per second	Creates per second
GemfireCluster|Region|GetsRate	Gets per second	Gets per second
GemfireCluster|Region|BucketCount	Bucket count	Bucket count
GemfireCluster|Region|AvgBucketSize	Average number of entries per bucket	Average number of entries per bucket
GemfireCluster|Region|Member|ActualRedundancy	Actual redundancy	Actual redundancy
GemfireCluster|Region|Member|BucketCount	Bucket count	Bucket count
GemfireCluster|Region|Member|AvgBucketSize	Average number of entries per bucket	Average number of entries per bucket
GemfireCluster|Region|Member|CreatesRate	Creates per second	Creates per second
GemfireCluster|Region|Member|GetsRate	Gets per second	Gets per second
GemfireCluster|Region|Member|DestroyRate	Destroys per second	Destroys per second
GemfireCluster|Region|Member|MissCount	Number of misses count	Number of cache misses
GemfireCluster|Region|Member|MissCountDelta	Number of new cache misses	Number of new cache misses

Table 7-79. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
GemfireCluster|Region|Member|HitCount	Number of hits count	Number of cache hits
GemfireCluster|Region|Member|HitCountDelta	Number of new cache hits	Number of new cache hits

Threshold Checking Metrics

Threshold checking metrics check the processed and computed metrics for the cluster.

Table 7-80. Threshold Checking Metrics for the Cluster

Metric Key	Metric Name	Description
ThresholdChecking|ProcessedMetricCount	Number of processed metrics	Number of processed metrics
ThresholdChecking|ProcessedMetricRate	Received metric processing rate (per second)	Received metric processing rate (per second)
ThresholdChecking|ComputedMetricCount	Number of computed metrics	Number of computed metrics
ThresholdChecking|ComputedMetricRate	Computed metric processing rate (per second)	Computed metric processing rate (per second)

Memory Metrics

Memory metrics provide memory CPU use information for the cluster.

Table 7-81. Memory Metrics for the Cluster

Metric Key	Metric Name	Description
Memory|AvgFreePhysicalMemory	Average free physical memory	Average free physical memory (GB)
Memory|TotalFreePhysicalMemory	Free physical memory	Free physical memory (GB)
Memory|TotalMemory	Total Available Memory	Total Available Memory (GB)
Memory|TotalUsedMemory	Actual Used Memory	Actual Used Memory (GB)
Memory|TotalDemandMemory	Memory Demand	Memory Demand (GB)

Elastic Memory Metrics

Elastic memory metrics provide reclaimable memory CPU use information for the cluster.

Table 7-82. Memory Metrics for the Cluster

Metric Key	Metric Name	Description
ElasticMemory|TotalMemory	Total Available Memory	Total Available Memory (GB)
ElasticMemory|TotalUsedMemory	Actual Used Memory	Actual Used Memory (GB)
ElasticMemory|TotalDemandMemory	Memory Demand	Memory Demand (GB)

CPU Metrics

CPU metrics provide CPU information for the cluster.

Table 7-83. CPU Metrics for the Cluster

Metric Key	Metric Name	Description
cpu|TotalCombinedUsage	CPU Load	CPU Load
cpu|TotalAvailable	CPU Available	CPU Available
cpu|TotalAvailable_ghz	Available	Available (GHz)
cpu|TotalUsage_ghz	Used	Used (GHz)
cpu|TotalUsage	CPU usage	CPU usage (%)

Disk Metrics

Disk metrics provide available disk information for the cluster.

Table 7-84. Disk Metrics for the Cluster

Metric Key	Metric Name	Description
Disk|DatabaseStorage|AvgAvailable	Average node disk available	Average node disk available
Disk|DatabaseStorage|MinAvailable	Minimum node disk available	Minimum node disk available
Disk|DatabaseStorage|MaxAvailable	Maximum node disk available	Maximum node disk available
Disk|DatabaseStorage|TotalAvailable	Available	Available
Disk|DatabaseStorage|Total	Total	Total
Disk|DatabaseStorage|TotalUsed	Used	Used
Disk|LogStorage|AvgAvailable	Average node disk available	Average node disk available
Disk|LogStorage|MinAvailable	Minimum node disk available	Minimum node disk available
Disk|LogStorage|MaxAvailable	Maximum node disk available	Maximum node disk available
Disk|LogStorage|TotalAvailable	Available	Available
Disk|LogStorage|Total	Total	Total
Disk|LogStorage|TotalUsed	Used	Used

Persistence Metrics

vRealize Operations Manager collects metrics for various persistence resources or service groups.

Activity Metrics

Activity metrics relate to the activity framework.

Table 7-85. Activity Metrics for Persistence

Metric Key	Metric Name	Description
Activity|RunningCount	Number Running	Number Running
Activity|ExecutedCount	Number Executed	Number Executed

Table 7-85. Activity Metrics for Persistence (continued)

Metric Key	Metric Name	Description
Activity|SucceededCount	Number Succeeded	Number Succeeded
Activity|FailedCount	Number Failed	Number Failed

Controller XDB Metrics

Controller metrics relate to the master database.

Table 7-86. Controller XDB Metrics for Persistence

Metric Key	Metric Name	Description
ControllerXDB|Size	Size	Size (Bytes)
ControllerXDB|TempDBSize	Temporary DB Size	Temporary DB Size (Bytes)
ControllerXDB|TotalObjectCount	Total Object Count	Total Object Count
ControllerXDB|AvgQueryDuration	Average Query Duration	Average Query Duration (ms)
ControllerXDB|MinQueryDuration	Minimum Query Duration	Minimum Query Duration (ms)
ControllerXDB|MaxQueryDuration	Maximum Query Duration	Maximum Query Duration (ms)
ControllerXDB|TotalTransactionCount	Total Transaction Count	Total Transaction Count
ControllerXDB|LockOperationErrorCount	Lock Operation Error Count	Lock Operation Error Count
ControllerXDB|DBCorruptionErrorCount	DB Corruption Error Count	DB Corruption Error Count
ControllerXDB| DBMaxSessionExceededCount	DB Maximum Sessions Exceeded Count	DB Maximum Sessions Exceeded Count
ControllerXDB|NumberWaitingForSession	Number of operations waiting for a session	Number of operations waiting for a session from the session pool
ControllerXDB|AvgWaitForSessionDuration	Average acquisition time from session pool	Average acquisition time from session pool
ControllerXDB|MinWaitForSessionDuration	Minimum acquisition time from session pool	Minimum acquisition time from session pool
ControllerXDB|MaxWaitForSessionDuration	Maximum acquisition time from session pool	Maximum acquisition time from session pool
ControllerXDB|TotalGetSessionCount	Total requests for a session from the session pool	Total requests for a session from the session pool
ControllerXDB|MaxActiveSessionCount	Maximum Concurrent Session Count	Maximum concurrent session count during the past collection interval.

Alarm SQL Metrics

Alarm metrics relate to the persistence of alerts and symptoms.

Table 7-87. Alarm XDB Metrics for Persistence

Metric Key	Metric Name	Description
AlarmSQL|Size	Size (Bytes)	Size (Bytes)
AlarmSQL|AvgQueryDuration	Average Query Duration (ms)	Average Query Duration (ms)
AlarmSQL|MinQueryDuration	Minimum Query Duration (ms)	Minimum Query Duration (ms)
AlarmSQL|MaxQueryDuration	Maximum Query Duration (ms)	Maximum Query Duration (ms)
AlarmSQL|TotalTransactionCount	Total Transaction Count	Total Transaction Count
AlarmSQL|TotalAlarms	Alarm Total Object Count	Alarm Total Object Count
AlarmSQL|TotalAlerts	Alert Total Object Count	Alert Total Object Count
AlarmSQL|AlertTableSize	Alert Table Size	Alert Table Size
AlarmSQL|AlarmTableSize	Alarm Table Size	Alarm Table Size

Key Value Store Database (KVDB)

KVDB metrics relate to the persistence of storing key-value data.

Metric Key	Metric Name	Description
KVDB|AvgQueryDuration	Average Query Duration	Average Query Duration
KVDB|MinQueryDuration	Minimum Query Duration	Minimum Query Duration
KVDB|MaxQueryDuration	Maximum Query Duration	Maximum Query Duration
KVDB|TotalTransactionCount	Total Transaction Count	Total Transaction Count

Historical Inventory Service XDB Metrics

Historical inventory service metrics relate to the persistence of configuration properties and their changes.

Table 7-88. Historical XDB Metrics for Persistence

Metric Key	Metric Name	Description
HisXDB|FunctionCalls|Count HisXDB|FunctionCalls	Number of Function calls	Number of Function calls
HisXDB|FunctionCalls|AvgDuration	Average execution time	Average execution time
HisXDB|FunctionCalls|MaxDuration	Max execution time	Max execution time
HisXDB|Size	Size	Size (Bytes)
HisXDB|TempDBSize	Temporary DB Size	Temporary DB Size (Bytes)
HisXDB|TotalObjectCount	Total Object Count	Total Object Count
HisXDB|AvgQueryDuration	Average Query Duration	Average Query Duration (ms)
HisXDB|MinQueryDuration	Minimum Query Duration	Minimum Query Duration (ms)
HisXDB|MaxQueryDuration	Maximum Query Duration	Maximum Query Duration (ms)
HisXDB|TotalTransactionCount	Total Transaction Count	Total Transaction Count

Table 7-88. Historical XDB Metrics for Persistence (continued)

Metric Key	Metric Name	Description
HisXDB|LockOperationErrorCount	Lock Operation Error Count	Lock Operation Error Count
HisXDB|DBCorruptionErrorCount	DB Corruption Error Count	DB Corruption Error Count
HisXDB|DBMaxSessionExceededCount	DB Maximum Sessions Exceeded Count	DB Maximum Sessions Exceeded Count
HisXDB|NumberWaitingForSession	Number of operations waiting for a session	Number of operations waiting for a session from the session pool
HisXDB|AvgWaitForSessionDuration	Average acquisition time from session pool	Average acquisition time from session pool
HisXDB|MinWaitForSessionDuration	Minimum acquisition time from session pool	Minimum acquisition time from session pool
HisXDB|MaxWaitForSessionDuration	Maximum acquisition time from session pool	Maximum acquisition time from session pool
HisXDB|TotalGetSessionCount	Total requests for a session from the session pool	Total requests for a session from the session pool
HisXDB|HisActivitySubmissionCount	HIS activity submission count	Number of Historical Inventory Service activities submitted
HisXDB|HisActivityCompletionCount	HIS activity completion count	Number of Historical Inventory Service activities completed
HisXDB|HisActivityCompletionDelayAvg	HIS activity average completion delay	The average amount of time from activity submission to completion
HisXDB|HisActivityCompletionDelayMax	HIS activity maximum completion delay	The maximum amount of time from activity submission to completion
HisXDB|HisActivityAbortedCount	HIS activity abort count	Number of Historical Inventory Service activities aborted

Remote Collector Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager remote collector node objects.

Table 7-89. Remote Collector Metrics

Metric Key	Metric Name	Description
ComponentCount	Component Count	The number of vRealize Operations Manager Objects reporting for this node.

Table 7-90. Memory Metrics for the Remote Collector

Metric Key	Metric Name	Description
mem|actualFree	Actual Free	Actual Free
mem|actualUsed	Actual Used	Actual Used

Table 7-90. Memory Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
mem|free	Free	Free)
mem|used	Used	Used
mem|total	Total	Total
mem|demand_gb	Estimated memory demand	Estimated memory demand

Table 7-91. Swap Metrics for the Remote Collector

Metric Key	Metric Name	Description
swap|total	Total	Total
swap|free	Free	Free
swap|used	Used	Used
swap|pageIn	Page in	Page in
swap|pageOut	Page out	Page out

Table 7-92. Resource limit Metrics for the Remote Collector

Metric Key	Metric Name	Description
resourceLimit|numProcesses	Number of processes	Number of processes
resourceLimit|openFiles	Number of open files	Number of open files
resourceLimit|openFilesMax	Number of open files maximum limit	Number of open files maximum limit
resourceLimit|numProcessesMax	Number of processes maximum limit	Number of processes maximum limit

Table 7-93. Network Metrics for the Remote Collector

Metric Key	Metric Name	Description
net|allInboundTotal	All inbound connections	All inbound total
net|allOutboundTotal	All outbound connections	All outbound total
net|tcpBound	TCP bound	TCP bound
net|tcpClose	TCP state CLOSE	Number of connections in TCP CLOSE
net|tcpCloseWait	TCP state CLOSE WAIT	Number of connections in TCP state CLOSE WAIT
net|tcpClosing	TCP state CLOSING	Number of connections in TCP state CLOSING
net|tcpEstablished	TCP state ESTABLISHED	Number of connections in TCP state ESTABLISHED
net|tcpIdle	TCP state IDLE	Number of connections in TCP state IDLE
net|tcpInboundTotal	TCP inbound connections	TCP inbound connections

Table 7-93. Network Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
net|tcpOutboundTotal	TCP outbound connections	TCP outbound connections
net|tcpLastAck	TCP state LAST ACK	Number of connections in TCP state LAST ACK
net|tcpListen	TCP state LISTEN	Number of connections in TCP state LISTEN
net|tcpSynRecv	TCP state SYN RCVD	Number of connections in TCP state SYN RCVD
net|tcpSynSent	TCP state SYN_SENT	Number of connections in TCP state SYN_SENT
net|tcpTimeWait	TCP state TIME WAIT	Number of connections in TCP state TIME WAIT

Table 7-94. Network Interface Metrics for the Remote Collector

Metric Key	Metric Name	Description
net|iface|speed	Speed	Speed (bits/sec)
net|iface|rxPackets	Receive packets	Number of received packets
net|iface|rxBytes	Receive bytes	Number of received bytes
net|iface|rxDropped	Receive packet drops	Number of received packets dropped
net|iface|rxFrame	Receive packets frame	Number of receive packets frame
net|iface|rxOverruns	Receive packets overruns	Number of receive packets overrun
net|iface|txPackets	Transmit packets	Number of transmit packets
net|iface|txBytes	Transmit bytes	Number of transmit bytes
net|iface|txDropped	Transmit packet drops	Number of transmit packets dropped
net|iface|txCarrier	Transmit carrier	Transmit carrier
net|iface|txCollisions	Transmit packet collisions	Number of transmit collisions
net|iface|txErrors	Transmit packet errors	Number of transmit errors
net|iface|txOverruns	Transmit packet overruns	Number of transmit overruns

Table 7-95. Disk Filesystem Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|fileSystem|total	Total	Total
disk|fileSystem|available	Available	Available
disk|fileSystem|used	Used	Used
disk|fileSystem|files	Total file nodes	Total number of file nodes
disk|fileSystem|filesFree	Total free file nodes	Total free file nodes

Table 7-95. Disk Filesystem Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
disk|fileSystem|queue	Disk queue	Disk queue
disk|fileSystem|readBytes	Read bytes	Number of bytes read
disk|fileSystem|writeBytes	Write bytes	Number of bytes written
disk|fileSystem|reads	Reads	Number of reads
disk|fileSystem|writes	Writes	Number of writes

Table 7-96. Disk Installation Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|installation|used	Used	Used
disk|installation|total	Total	Total
disk|installation|available	Available	Available

Table 7-97. Disk Database Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|db|used	Used	Used
disk|db|total	Total	Total
disk|db|available	Available	Available

Table 7-98. Disk Log Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|log|used	Used	Used
disk|log|total	Total	Total
disk|log|available	Available	Available

Table 7-99. CPU Metrics for the Remote Collector

Metric Key	Metric Name	Description
cpu|combined	Combined load	Combined load (User + Sys + Nice + Wait)
cpu|idle	Idle	Idle time fraction of total available cpu (cpu load)
cpu|irq	Irq	Interrupt time fraction of total available cpu (cpu load)
cpu|nice	Nice	Nice time fraction of total available cpu (cpu load)
cpu|softIrq	Soft Irq	Soft interrupt time fraction of total available cpu (cpu load)
cpu|stolen	Stolen	Stolen time fraction of total available cpu (cpu load)

Table 7-99. CPU Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
cpu|sys	Sys	Sys time fraction of total available cpu (cpu load)
cpu|user	User	User time fraction of total available cpu (cpu load)
cpu|wait	Wait	Wait time fraction of total available cpu (cpu load)
cpu|total	Total available for a cpu	Total available for a cpu
cpu|allCpuCombined	Total combined load for all cpus	Total combined load for all cpus (cpu load)
cpu|allCpuTotal_ghz	Available	Available
cpu|allCpuCombined_ghz	Used	Used
cpu|allCpuCombined_percent	CPU usage	CPU usage (%)

Table 7-100. Device Metrics for the Remote Collector

Metric Key	Metric Name	Description
device|iops	Reads/writes per second	Average number of read/write commands issued per second during the collection interval
device|await	Average transaction time	Average transaction time (milliseconds)

Table 7-101. Service Metrics for the Remote Collector

Metric Key	Metric Name	Description
service|proc|fdUsage	Total number of open file descriptors	Total number of open file descriptors (Linux). Total number of open handles (Windows)

Table 7-102. NTP Metrics for the Remote Collector

Metric Key	Metric Name	Description
ntp|serverCount	Configured server count	Configured server count
ntp|unreachableCount	Unreachable server count	Unreachable server count
ntp|unreachable	Unreachable	Is the NTP server unreachable. Value of 0 is reachable, 1 means the server was not reached or didn't respond.

Metrics for vRealize Automation

The vRealize Automation solution collects metrics about the total number of virtual machines deployed.

Table 7-103. Metrics

Metric Name	Description
Total Deployed	The total number of virtual machines deployed from the blueprint. Key: DeploymentCount | TotalDeployed

Managed Resources Object as a Filter in vRealize Automation

The vRealize Automation solution uses filters to display the VMware vCenter adapter objects that are managed by or have some association with vRealize Automation. Some of the dashboards have widgets that are configured to display only VMware vCenter adapter objects that vRealize Automation manages or is associated with. vRealize Automation uses an object called Managed Resources as a filter to display only those objects. All of these resources are placed under the Managed Resources object of type vRealize Automation Entity Status. The absence of this filter causes all VMware vCenter adapter objects to be displayed in the widgets. If you delete the Managed Resources object, the adapter re-creates the object, but the dashboards display incorrect information in the widgets that use this filter. If you delete the Managed Resources object, you must manually configure the widgets in the dashboard and select the Managed Resources object as a filter in each section that displays the VMware vCenter adapter objects.

Metrics for vSAN

vRealize Operations Manager collects metrics for vSAN objects.

In the menu, click **Environment > All Objects > vSAN Adapter**. Select one of the vSAN adapter objects listed and click the **All Metrics** tab.

Disk I/O and Disk Space Metrics for vSAN Disk Groups

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN disk groups.

Disk I/O metrics for the vSAN disk groups include:

- Disk I/O|Reads Per Second (IOPS)
- Disk I/O|Writes Per Second (IOPS
- Disk I/O|Max Observed Reads Per Second (IOPS)
- Disk I/O|Max Observed Writes Per Second (IOPS)
- Disk I/O|Throughput Read (bps)
- Disk I/O|Throughput Write (bps)
- Disk I/O|Average Read Latency (ms)
- Disk I/O|Average Write Latency (ms)
- Disk I/O|Read Count
- Disk I/O|Write Count

- Disk I/O|Average Device Latency
- Disk I/O|Average Device Read Latency
- Disk I/O|Average Device Write Latency
- Disk I/O|Total Bus Resets
- Disk I/O|Total Commands Aborted per second
- Disk I/O|Total Number of Errors

Disk space metrics for vSAN disk groups include:

- Disk Space|Capacity (bytes)
- Disk Space|Used (bytes)
- Disk Space|Usage (%)

Read Cache Metrics for vSAN Disk Groups

The vRealize Operations Manager collects metrics and performs capacity trend analysis on a hybrid vSAN read cache. Read Cache metrics are not collected for a vSAN all-flash configuration.

Read cache metrics for the vSAN disk group include:

- Read Cache|Hit Rate (%)
- Read Cache|Miss Rate Ratio
- Read Cache|Read Cache Reads Per Second (IOPS)
- Read Cache|Read Cache Read Latency (ms)
- Read Cache|Read Cache Read I/O Count
- Read Cache|Read Cache Writes Per Second (IOPS)
- Read Cache|Read Cache Write Latency (ms)
- Read Cache|Read Cache Write I/O Count

Write Buffer Metrics for vSAN Disk Groups

The vRealize Operations Manager collects the metrics you use to monitor the write buffer capacity of your vSAN disk groups.

A reasonably balanced system consumes a significant amount of write buffer. Before placing additional workload on the vSAN, check the write buffer metrics for the vSAN disk group.

- Write Buffer|Capacity (bytes)
- Write Buffer|Free (%)
- Write Buffer|Usage (%)
- Write Buffer|Used (byte)
- Write Buffer|Write Buffer Reads Per Second (IOPS)

- Write Buffer|Write Buffer Read Latency (ms)
- Write Buffer|Write Buffer Read I/O Count
- Write Buffer|Write Buffer Writes Per Second (IOPS)
- Write Buffer|Write Buffer Write Latency (ms)
- Write Buffer|Write Buffer Write I/O Count

Congestion Metrics for vSAN Disk Groups

The vRealize Operations Manager collects congestion metrics for the vSAN disk group.

- Congestion| Memory Congestion - Favorite
- Congestion| SSD Congestion - Favorite
- Congestion| IOPS Congestion - Favorite
- Congestion| Slab Congestion
- Congestion| Log Congestion
- Congestion| Comp Congestion

Metrics for vSAN Cluster

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN cluster.

Metrics for vSAN cluster include:

Component	Metrics
Component Limit	■ vSAN|Component Limit|Component Limit Used (%) ■ vSAN|Component Limit|Total Component Limit ■ vSAN|Component Limit|Used Component Limit
Disk Space	■ vSAN|Disk Space|Disk Space Used (%) ■ vSAN|Disk Space|Total Disk Space (GB) ■ vSAN|Disk Space|Used Disk Space (GB)
Read Cache	■ vSAN|Read Cache|Read Cache Reserved (%) ■ vSAN|Read Cache|Reserved Read Cache Size (GB) ■ vSAN|Read Cache|Total Read Cache Size (GB)

Component	Metrics
Performance	■ vSAN|Read Cache|Reads Per Second (IOPS) ■ vSAN|Read Cache|Read Throughput (KBps) ■ vSAN|Read Cache|Average Read Latency (ms) ■ vSAN|Read Cache|Writes Per Second (IOPS) ■ vSAN|Read Cache|Write Throughput (KBps) ■ vSAN|Read Cache|Average Write Latency (ms) ■ vSAN|Read Cache|Congestion ■ vSAN|Read Cache|Outstanding I/O ■ vSAN|Read Cache|Total IOPS ■ vSAN|Read Cache|Total Latency (ms) ■ vSAN|Read Cache|Total Throughput (KBps)
Deduplication And Compression Overview	■ vSAN|Deduplication And Compression Overview|Used Before ■ vSAN|Deduplication And Compression Overview|Used After ■ vSAN|Deduplication And Compression Overview|Savings ■ vSAN|Deduplication And Compression Overview|Ratio
Summary	■ Summary|Number of Cache Disks ■ Summary|Total Number of Capacity Disks ■ Summary|CPU Workload ■ Summary|Memory Workload ■ Summary|Total Number of Disk Groups ■ Summary|Total Active Alerts Count ■ Summary|Total Number of VMs ■ Summary|Total Number of Hosts ■ Summary|vSAN Cluster Capacity Remaining (%) ■ Summary|VSAN Cluster Storage Time Remaining ■ Summary|vSAN Capacity Disk Used

Metrics for vSAN Enabled Host

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN enabled host.

Metrics for vSAN enabled host include:

Component	Metrics
Component Limit	■ vSAN|Component Limit|Component Limit Used (%) ■ vSAN|Component Limit|Total Component Limit ■ vSAN|Component Limit|Used Component Limit
Disk Space	■ vSAN|Disk Space|Disk Space Used (%) ■ vSAN|Disk Space|Total Disk Space (GB) ■ vSAN|Disk Space|Used Disk Space (GB)
Read Cache	■ vSAN|Read Cache|Read Cache Reserved (%) ■ vSAN|Read Cache|Reserved Read Cache Size (GB) ■ vSAN|Read Cache|Total Read Cache Size (GB)

Metrics for vSAN Datastore

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN datastore.

Datastore I/O metrics for vSAN datastore include:

- Datastore I/O|Reads Per Second (IOPS)
- Datastore I/O|Read Rate (KBps)
- Datastore I/O|Read Latency (ms)
- Datastore I/O|Writes Per Second (IOPS)
- Datastore I/O|Write Rate (KBps)
- Datastore I/O|Write Latency (ms)
- Datastore I/O|Outstanding I/O requests
- Datastore I/O|Congestion

Metrics for vSAN Cache Disk

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN cache disk.

Metrics for vSAN cache disk include:

Component	Metrics
Performance	■ Performance|Bus Resets ■ Performance|Commands Aborted Per Second ■ Performance|Device Latency (ms) ■ Performance|Device Read Latency (ms) ■ Performance|Device Write Latency (ms) ■ Performance|Read Requests Per Second ■ Performance|Average Reads Per Second ■ Performance|Write Requests Per Second ■ Performance|Average Writes Per Second ■ Performance|Read Rate ■ Performance|Write Rate ■ Performance|Usage ■ Performance|HDD Errors
SCSI SMART Statistics	■ SCSI SMART Statistics|Health Status ■ SCSI SMART Statistics|Media Wearout Indicator ■ SCSI SMART Statistics|Write Error Count ■ SCSI SMART Statistics|Read Error Count ■ SCSI SMART Statistics|Power on Hours ■ SCSI SMART Statistics|Reallocated Sector Count ■ SCSI SMART Statistics|Raw Read Error Rate ■ SCSI SMART Statistics|Drive Temperature ■ SCSI SMART Statistics|Maximum Observed Drive Temperature ■ SCSI SMART Statistics|Drive Rated Max Temperature ■ SCSI SMART Statistics|Write Sectors TOT Count ■ SCSI SMART Statistics|Read Sectors TOT Count ■ SCSI SMART Statistics|Initial Bad Block Count
Capacity	■ vSAN|Health|Capacity|Total Disk Capacity (GB) ■ vSAN|Health|Capacity|Used Disk Capacity (GB)
Congestion Health	■ vSAN|Health|Congestion Health|Congestion Value
Performance	■ vSAN|Performance|Physical Layer Reads Per Second ■ vSAN|Performance|Physical Layer Writes Per Second ■ vSAN|Performance|Physical Layer Read Throughput (KBps) ■ vSAN|Performance|Physical Layer Write Throughput (KBps) ■ vSAN|Performance|Physical Layer Read Latency (ms) ■ vSAN|Performance|Physical Layer Write Latency (ms) ■ vSAN|Performance|Physical Layer Read Count ■ vSAN|Performance|Physical Layer Write Count ■ vSAN|Performance|Device Average Latency (ms) ■ vSAN|Performance|Guest Average Latency (ms)

Metrics for vSAN Capacity Disk

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN capacity disk.

Metrics for vSAN capacity disk include:

Component	Metrics
Performance	■ Performance|Bus Resets ■ Performance|Commands Aborted Per Second ■ Performance|Device Latency (ms) ■ Performance|Device Read Latency (ms) ■ Performance|Device Write Latency (ms) ■ Performance|Read Requests Per Second ■ Performance|Average Reads Per Second ■ Performance|Write Requests Per Second ■ Performance|Average Writes Per Second ■ Performance|Read Rate ■ Performance|Write Rate ■ Performance|Usage ■ Performance|HDD Errors
SCSI SMART Statistics	■ SCSI SMART Statistics|Health Status ■ SCSI SMART Statistics|Media Wearout Indicator ■ SCSI SMART Statistics|Write Error Count ■ SCSI SMART Statistics|Read Error Count ■ SCSI SMART Statistics|Power on Hours ■ SCSI SMART Statistics|Reallocated Sector Count ■ SCSI SMART Statistics|Raw Read Error Rate ■ SCSI SMART Statistics|Drive Temperature ■ SCSI SMART Statistics|Maximum Observed Drive Temperature ■ SCSI SMART Statistics|Drive Rated Max Temperature ■ SCSI SMART Statistics|Write Sectors TOT Count ■ SCSI SMART Statistics|Read Sectors TOT Count ■ SCSI SMART Statistics|Initial Bad Block Count
Capacity	■ vSAN|Health|Total Disk Capacity (GB) ■ vSAN|Health|Used Disk Capacity (GB)

Component	Metrics
Congestion Health	vSAN|Health|Congestion Value
Performance	■ vSAN|Performance|Physical Layer Reads Per Second ■ vSAN|Performance|Physical Layer Writes Per Second ■ vSAN|Performance|Physical Layer Read Throughput (KBps) ■ vSAN|Performance|Physical Layer Write Throughput (KBps) ■ vSAN|Performance|Physical Layer Read Latency (ms) ■ vSAN|Performance|Physical Layer Write Latency (ms) ■ vSAN|Performance|Physical Layer Read Count ■ vSAN|Performance|Physical Layer Write Count ■ vSAN|Performance|Device Average Latency (ms) ■ vSAN|Performance|Guest Average Latency (ms) ■ vSAN|Performance|vSAN Layer Reads Per Second ■ vSAN|Performance|vSAN Layer Writes Per Second ■ vSAN|Performance|vSAN Layer Read Latency (ms) ■ vSAN|Performance|vSAN Layer Write Latency (ms) ■ vSAN|Performance|vSAN Layer Read Count ■ vSAN|Performance|vSAN Layer Write Count

Properties for vSAN capacity disk include:

- Name
- Size
- Vendor
- Type
- Queue Depth

Metrics for vSAN World

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN world.

Metrics for vSAN world include:

- Summary|Total Number of VMs
- Summary|Total Number of Hosts
- Summary|Total IOPS
- Summary|Total Latency
- Summary|Total Number of Clusters
- Summary|Total Number of DiskGroups
- Summary|Total Number of Cache Disks
- Summary|Total Number of Capacity Disks

- Summary|Total Number of Datastores
- Summary|Total vSAN Disk Capacity (TB)
- Summary|Total vSAN Disk Capacity Used (TB)
- Summary|Remaining Capacity (TB)
- Summary|Remaining Capacity (%)
- Summary|Total Savings by Deduplication and Compression (GB)

Metrics for the Operating Systems and Remote Service Monitoring Plug-ins in End Point Operations Management

vRealize Operations Manager collects metrics for the object types in the Operating Systems and Remote Service Monitoring plug-ins.

Due to rounding in metric time calculation, there can be situations in which the Resource Availability metric is rounded up. Rounding up the metric appears as gaps in the metrics reported by the End Point Operations Management agent. However, the metrics are fully reported.

Operating Systems Plug-in Metrics

The Operating Systems plug-in collects metrics for object types such Linux, AIX, Solaris, and Windows. The Operating Systems plug-in also collects metrics for Windows services, Script services, and Multiprocess services.

AIX Metrics

The Operating Systems Plug-in discovers the metrics for the AIX object type. AIX 6.1 and 7.1 are supported.

Table 7-104. AIX metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	True
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False

Table 7-104. AIX metrics (continued)

Name	Category	KPI
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Cpu Wait Time	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False

Table 7-104. AIX metrics (continued)

Name	Category	KPI
Nfs Server V3 Read per Minute	UTILIZATION	False
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Symlink	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False
Nfs Server V3 Readdir per Minute	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False

Table 7-104. AIX metrics (continued)

Name	Category	KPI
Page Major faults	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Page Faults	UTILIZATION	False
Percent Used Swap	UTILIZATION	True
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Pages In	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Linux Metrics

The Operating Systems Plug-in discovers the metrics for the Linux object type.

Table 7-105. Linux Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False

Table 7-105. Linux Metrics (continued)

Name	Category	KPI
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Stolen	UTILIZATION	False
Cpu Wait Time	UTILIZATION	False
Cpu Irq Time per Minute	UTILIZATION	False
Cpu Softlrq Time	UTILIZATION	False
Cpu Stolen Time per Minute	UTILIZATION	False
Cpu Stolen Time	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Irq	UTILIZATION	False
Cpu Softlrq Time per Minute	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Irq Time	UTILIZATION	False
Cpu Softlrq	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Free Memory (+ buffers/cache)	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False

Table 7-105. Linux Metrics (continued)

Name	Category	KPI
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False
Nfs Server V3 Read per Minute	UTILIZATION	False
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False
Nfs Server V3 Readdir per Minute	UTILIZATION	False

Table 7-105. Linux Metrics (continued)

Name	Category	KPI
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Page Major faults	UTILIZATION	False
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Page Faults	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
Used Memory (- buffers/cache)	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False

Table 7-105. Linux Metrics (continued)

Name	Category	KPI
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Solaris Metrics

The Operating Systems Plug-in discovers the metrics for the Solaris object type. Solaris x86 and SPARC are supported.

Table 7-106. Solaris Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False

Table 7-106. Solaris Metrics (continued)

Name	Category	KPI
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Wait Time	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False
Nfs Server V3 Read per Minute	UTILIZATION	False

Table 7-106. Solaris Metrics (continued)

Name	Category	KPI
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Symlink	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False
Nfs Server V3 Readdir per Minute	UTILIZATION	False
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Page Major faults	UTILIZATION	False

Table 7-106. Solaris Metrics (continued)

Name	Category	KPI
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Page Faults	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Microsoft Windows Metrics

The Operating Systems Plug-in discovers the metrics for the Microsoft Windows object type. Microsoft Windows Server 2012 R2 and 2008 R2 are supported.

Table 7-107. Microsoft Windows Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
Avg. Disk sec/Transfer	THROUGHPUT	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False

Table 7-107. Microsoft Windows Metrics (continued)

Name	Category	KPI
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Free Memory	UTILIZATION	False
Memory Page Faults/sec	UTILIZATION	False
Memory System Driver Resident Bytes	UTILIZATION	False
Memory Available Bytes	UTILIZATION	False
Memory System Driver Total Bytes	UTILIZATION	False
Memory % Committed Bytes In Use	UTILIZATION	False
Memory Standby Cache Core Bytes	UTILIZATION	False
Memory Transition Pages RePurposed/sec	UTILIZATION	False
Memory Write Copies/sec	UTILIZATION	False
Memory Available KBytes	UTILIZATION	False
Memory Page Reads/sec	UTILIZATION	False
Memory Committed Bytes	UTILIZATION	False
Memory Pool Nonpaged Bytes	UTILIZATION	False
Memory System Code Resident Bytes	UTILIZATION	False
Memory Page Writes/sec	UTILIZATION	False
Memory Available MBytes	UTILIZATION	False
Memory Standby Cache Normal Priority Bytes	UTILIZATION	False
Memory Pages/sec	UTILIZATION	False
Memory Modified Page List Bytes	UTILIZATION	False
Memory Cache Faults/sec	UTILIZATION	False
Memory Pool Nonpaged Allocs	UTILIZATION	False
Memory System Code Total Bytes	UTILIZATION	False
Memory Pool Paged Allocs	UTILIZATION	False
Memory Pages Input/sec	UTILIZATION	False
Memory Pool Paged Bytes	UTILIZATION	False
Memory Pool Paged Resident Bytes	UTILIZATION	False
Memory Cache Bytes	UTILIZATION	False
Memory Standby Cache Reserve Bytes	UTILIZATION	False

Table 7-107. Microsoft Windows Metrics (continued)

Name	Category	KPI
MemoryFreeSystemPageTableEntries	UTILIZATION	False
Memory Free %26 Zero Page List Bytes	UTILIZATION	False
Memory System Cache Resident Bytes	UTILIZATION	False
Memory Cache Bytes Peak	UTILIZATION	False
Memory Commit Limit	UTILIZATION	False
Memory Transition Faults/sec	UTILIZATION	False
Memory Pages Output/sec	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	True
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Windows Service Metrics

The Operating Systems Plug-in discovers the metrics for Windows service.

Table 7-108. Windows Services Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Start Time	AVAILABILITY	False
Start Type	AVAILABILITY	False
Cpu User Time	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Total Time per Minute	UTILIZATION	False
Cpu System Time per Minute	UTILIZATION	False
Cpu Total Time	UTILIZATION	False
Cpu User Time per Minute	UTILIZATION	False
Cpu System Time	UTILIZATION	False
Memory Size	UTILIZATION	True
Open Handles	UTILIZATION	False
Resident Memory Size	UTILIZATION	False
Threads	UTILIZATION	False

If you stop an End Point Operations Management agent by using Windows Services, and remove the `data` directory from inside the agent installation directory, when you start the agent again, using Windows Services, no metrics are collected. 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`.

Script Metrics

The Operating Systems Plug-in discovers the metrics for the Script service.

Table 7-109. Script Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Execution Time	THROUGHPUT	True
Result Value	UTILIZATION	True

Multiprocess Service Metrics

The Operating Systems Plug-in discovers the metrics for the Multiprocess service.

Table 7-110. Multiprocess Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Cpu User Time	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Total Time per Minute	UTILIZATION	False
Cpu System Time per Minute	UTILIZATION	False
Cpu Total Time	UTILIZATION	False
Cpu User Time per Minute	UTILIZATION	False
Cpu System Time	UTILIZATION	False
Memory Size	UTILIZATION	True
Number of Processes	UTILIZATION	False
Resident Memory Size	UTILIZATION	False

NFS Metrics

The End Point Operations Management agents collect metrics for the NFS-mounted file systems.

The following metrics are collected.

Name	Category
Resource Availability	Availability
Use Percent (%)	Utilization
Total Bytes Free (KB)	Utilization

Remote Service Monitoring Plug-in Metrics

The Remote Service Monitoring plug-in collects metrics for object types such HTTP Check, TCP Check, and ICMP Check.

HTTP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the HTTP Check object type.

Table 7-111. HTTP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Last Modified	AVAILABILITY	False
State CLOSE	THROUGHPUT	False
State CLOSE_WAIT	THROUGHPUT	False
State ESTABLISHED	THROUGHPUT	False
Inbound Connections	THROUGHPUT	False
State TIME_WAIT	THROUGHPUT	False

Table 7-111. HTTP Check Metrics (continued)

Name	Category	KPI
All Inbound Connections	THROUGHPUT	False
State SYN_SENT	THROUGHPUT	False
State FIN_WAIT2	THROUGHPUT	False
Outbound Connections	THROUGHPUT	False
State LAST_ACK	THROUGHPUT	False
Response Time	THROUGHPUT	True
State CLOSING	THROUGHPUT	False
All Outbound Connections	THROUGHPUT	False
State SYN_RECV	THROUGHPUT	False
State FIN_WAIT1	THROUGHPUT	False
Response Code	UTILIZATION	True

ICMP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the ICMP Check object type.

Table 7-112. ICMP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Response Time	THROUGHPUT	True

TCP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the TCP Check object type.

Table 7-113. TCP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Response Time	THROUGHPUT	True
State CLOSE	THROUGHPUT	False
State CLOSE_WAIT	THROUGHPUT	False
State ESTABLISHED	THROUGHPUT	False
Inbound Connections	THROUGHPUT	False
State TIME_WAIT	THROUGHPUT	False
All Inbound Connections	THROUGHPUT	False
State SYN_SENT	THROUGHPUT	False
State FIN_WAIT2	THROUGHPUT	False
Outbound Connections	THROUGHPUT	False

Table 7-113. TCP Check Metrics (continued)

Name	Category	KPI
State LAST_ACK	THROUGHPUT	False
State CLOSING	THROUGHPUT	False
All Outbound Connections	THROUGHPUT	False
State SYN_RECV	THROUGHPUT	False
State FIN_WAIT1	THROUGHPUT	False

Alert Definitions in vRealize Operations Manager

Alert definitions are a combination of symptoms and recommendations that identify problem areas in vRealize Operations Manager and generate alerts on which you act for those areas.

Alert definitions are provided for various objects in your environment. You can also create your own alert definitions. See Create an Alert Definition for Department Objects.

- Cluster Compute Resource Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the Cluster Compute Resource objects in your environment.

- Host System Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the Host System objects in your environment.

- vRealize Automation 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.

- vSAN Alerts Definitions

  vRealize Operations Manager generates an alert if a problem occurs with the components in the storage area network that the vSAN adapter is monitoring.

- Alerts in the vSphere Web Client

  The vSphere Web Client displays the results of health tests for the following vSAN monitored groups:

- vSphere Distributed Port Group

  The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Port objects in your environment.

- Virtual Machine Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the virtual machine objects in your environment.

- vSphere Distributed Switch Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Switch objects in your environment.

- vCenter Server Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the vCenter Server objects in your environment.

- Datastore Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

- Data Center Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the Data Center objects in your environment.

- Custom Data Center Alert Definitions

  The vCenter adapter provides alert definitions that generate alerts on the Custom Data Center objects in your environment.

Cluster Compute Resource Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Cluster Compute Resource objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has CPU contention caused by less than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ > 0 descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/ critical level] ■ <= 50% of descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 Use the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has CPU contention caused by more than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ Cluster CPU demand at warning/ immediate/critical level ■ > 50% of descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase CPU capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has CPU contention caused by overpopulation of virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ Cluster CPU workload at warning/ immediate/critical level ■ = 0 descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/ critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase CPU capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has unexpected high CPU workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU workload above DT ■ Cluster CPU workload at warning/ immediate/critical level	1 Check the applications running on the virtual machines in the cluster to determine whether high CPU workload is an expected behavior. 2 Add more hosts to the cluster to increase CPU capacity. 3 Use vSphere vMotion to migrate some virtual machines to a different cluster if possible.
Fully-automated DRS-enabled cluster has memory contention caused by less than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ > 0 descendant virtual machines have [Virtual machine memory workload at warning /immediate/ critical level] ■ <= 50% of descendant virtual machines have [Virtual machine memory workload at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has memory contention caused by more than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level ■ > 50% of descendant virtual machines have [Virtual machine memory demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has memory contention caused by overpopulation of virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level ■ = 0 descendant virtual machines have [Virtual machine memory demand at warning /immediate/ critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
More than 5% of virtual machines in the cluster have memory contention due to memory compression, ballooning or swapping.	■ Virtual machine memory limit is set AND ■ > 5% of descendant virtual machines have [virtual machine memory contention is at warning/ immediate/critical level] AND ■ > 5% of descendant virtual machines have [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/ immediate/critical level]	1 Add more hosts to the cluster to increase memory capacity. 2 Use vMotion to migrate some virtual machines off the host or cluster.
Fully-automated DRS-enabled cluster has unexpected high memory workload and contention.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention above DT ■ Cluster memory content is at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level	1 Check the applications running on the virtual machines in the cluster to determine whether high memory workload is an expected behavior. 2 Add more hosts to the cluster to increase memory capacity. 3 Use vSphere vMotion to migrate some virtual machines to a different cluster if possible.

Alert Definition	Symptoms	Recommendations
vSphere HA failover resources are insufficient.	vSphere HA failover resources are insufficient (fault symptom)	To resolve this problem, use similar CPU and memory reservations for all virtual machines in the cluster. If this solution is not possible, consider using a different vSphere HA admission control policy, such as reserving a percentage of cluster resource for failover. Alternatively, you can use advanced options to specify a cap for the slot size. For more information, see the vSphere Availability Guide. Hosts that have vSphere HA agent errors are not good candidates for providing failover capacity in the cluster and their resources are not considered for vSphere HA admission control purposes. If many hosts have a vSphere HA agent error, vCenter Server generates this event leading to the fault. To resolve vSphere HA agent errors, check the event logs for the hosts to determine the cause of the errors. After you resolve any configuration problems, reconfigure vSphere HA on the affected hosts or on the cluster.
vSphere HA master missing.	vCenter Server is unable to find a master vSphere HA agent (fault symptom)	Check the fault page under the **Analysis** tab for this object to find more objects.

Host System Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Host System objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Standalone host has CPU contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] ■ <= 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/ critical level]	Use 1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has CPU contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ Host CPU demand at warning/ immediate/critical level ■ > 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/ critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has CPU contention due to overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ Host CPU demand at warning/ immediate/critical level ■ = 0 child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Host in a cluster that does not have fully-automated DRS enabled has contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention is at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] ■ <= 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/ critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention at warning/ immediate/critical level ■ Host CPU demand at warning/ immediate/critical level ■ > 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/ critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention at warning/ immediate/critical level ■ Host CPU demand at warning/ immediate/critical level ■ = 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Standalone host has memory contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has memory contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Standalone host has memory contention due to overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/ critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by less than half of the virtual machines.	Symptoms include the following: ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory contention at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/ critical level] ■ <= 50% of child virtual machines have [Virtual machine memory workload at warning/ immediate/ critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host is experiencing high number of received or transmitted packets dropped.	Symptoms include the following: ■ Host network received packets dropped ■ Host network transmitted packets dropped	1 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic. 2 Verify the health of the physical network adapter, configuration, driver and firmware versions. 3 Contact VMware support.
Host is experiencing high number of received packets dropped.	Symptoms include the following: ■ Host network received packets dropped ■ Host network received packets dropped above DT ■ Host network data receive workload at Warning level ■ Host network data receive workload above DT ■ Host CPU demand at Critical level	1 If the host has one CPU, upgrade the host or use a host that has larger CPU capacity. 2 Add an additional NIC to the host. 3 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.
Host is experiencing high number of transmitted packets dropped.	Symptoms include the following: ■ Host network transmitted packets dropped ■ Host network transmitted packets dropped above DT ■ Host network data transmit workload at Warning level ■ Host network data transmit workload above DT ■ Host is dropping high percentage of packets	1 Add an additional NIC to the host. 2 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.

Alert Definition	Symptoms	Recommendations
ESXi host has detected a link status 'flapping' on a physical NIC.	Physical NIC link state flapping (fault symptom).	ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert.
ESXi host has detected a link status down on a physical NIC.	Physical NIC link state down (fault symptom).	ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert.
Battery sensors are reporting problems.	Symptoms include the following: ■ Battery sensor health is red OR ■ Battery sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Baseboard Management Controller sensors are reporting problems.	Symptoms include the following: ■ Baseboard Management Controller sensor health is red OR ■ Baseboard Management Controller sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Fan sensors are reporting problems.	■ Fan sensor health is red OR ■ Fan sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Hardware sensors are reporting problems.	■ Hardware sensor health is red OR ■ Hardware sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Memory sensors are reporting problems.	■ Memory sensor health is red OR ■ Memory sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.

Alert Definition	Symptoms	Recommendations
Power sensors are reporting problems.	■ Power sensor health is red OR ■ Power sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Processor sensors are reporting problems.	■ Processor sensor health is red ■ Processor sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
SEL sensors are reporting problems.	■ SEL sensor health is red OR ■ SEL sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Storage sensors are reporting problems.	■ Storage sensor health is red OR ■ Storage sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
System Board sensors are reporting problems.	■ System board sensor health is red OR ■ System board sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Temperature sensors are reporting problems.	■ Temperature sensor health is red OR ■ Temperature sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Voltage sensors are reporting problems.	■ Voltage sensor health is red OR ■ Voltage sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Critical

Alert Definition	Symptoms	Recommendations
Host has lost connection to vCenter.	Host disconnected from vCenter	Click "Open Host in vSphere Web Client" in the Actions menu at the top of Alert details page to connect to the vCenter managing this host and manually reconnect the host to vCenter Server. After the connection to the host is restored by vCenter Server, the alert will be canceled.
vSphere High Availability (HA) has detected a network-isolated host.	vSphere HA detected a network isolated host (fault symptom).	Resolve the networking problem that prevents the host from pinging its isolation addresses and communicating with other hosts. Make sure that the management networks that vSphere HA uses include redundancy. With redundancy, vSphere HA can communicate over more than one path, which reduces the chance of a host becoming isolated.
vSphere High Availability (HA) has detected a possible host failure.	vSphere HA detected a host failure (fault symptom).	Find the computer that has the duplicate IP address and reconfigure it to have a different IP address. This fault is cleared and the alert canceled when the underlying problem is resolved, and the vSphere HA master agent is able to connect to the HA agent on the host. **Note** You can use the Duplicate IP warning in the `/var/log/vmkernel` log file on an ESX host or the /var/log/messages log file on an ESXi host to identify the computer that has the duplicate IP address.

Alert Definition	Symptoms	Recommendations
Host is experiencing network contention caused by too much traffic.	Symptoms include all the following: ■ Host is experiencing dropped network packets ■ Host network workload at warning/immediate/critical level	1 Review the load balancing policy in the Port Group and the vSwitch. 2 Add an additional NIC to the host. 3 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.
The host has lost connectivity to a dvPort.	Lost network connectivity to dvPorts (fault symptom).	Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the dvPort.

Alert Definition	Symptoms	Recommendations
The host has lost connectivity to the physical network.	Lost network connectivity (fault symptom).	To determine the actual failure or to eliminate possible problems, check the status of the vmnic in the vSphere Client or from the ESX service console: ■ To check the status in the vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently down. ■ From the service console, run the command:`esxcfg-nics`. The output that appears is similar to the following: Name PCI Driver Link Speed Duplex Description ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, bring down each ESX host port on the physical switch, run `esxcfg-nics -l"`, and observe the affected vmnics. Verify that the vmnic identified in the alert is still connected to the switch and configured properly: ■ Make sure that the network cable is still connected to the switch and to the host. ■ Make sure that the switch is connected to the system, is still functioning properly, and has not been inadvertently misconfigured. For more information, see the switch documentation. ■ Check for activity between the physical switch and the vmnic. You can check activity by performing a network trace or observing activity LEDs.

Alert Definition	Symptoms	Recommendations
		■ Check for network port settings on the physical switch. To reconfigure the service console IP address if the affected vmnic is associated with a service console, see http://kb.vmware.com/kb/1000258 If the problem is caused by your hardware, contact your hardware vendor for replacement hardware.
The host lost connectivity to a Network File System (NFS) server.	Lost connection to NFS server (fault symptom).	1 Verify the NFS server is running. 2 Check the network connection to make sure the ESX host can connect to the NFS server. 3 Determine whether the other hosts that use the same NFS mount are experiencing the same problem, and check the NFS server status and share points. 4 Make sure that you can reach the NFS server by logging into the service console and using `vmkping` to ping the NFS server: "vmkping <nfs server>". 5 For advanced troubleshooting information, seehttp://kb.vmware.com/kb/1003967
A fatal error occurred on a PCIe bus during system reboot.	A fatal PCIe error occurred.	Check and replace the PCIe device identified in the alert as the cause of the problem. Contact the vendor for assistance.
A fatal memory error was detected at system boot time.	A fatal memory error occurred.	Replace the faulty memory or contact the vendor.

Health/Immediate

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Immediate

Alert Definition	Symptom	Recommendations
The host has lost redundant connectivity to a dvPort.	Lost network redundancy to DVPorts (fault symptom).	Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the DVPort.
The host has lost redundant uplinks to the network.	Lost network redundancy (fault symptom).	To determine the actual failure or to eliminate possible problems, first connect to ESX through SSH or the console: 1 Identify the available uplinks by running `esxcfg-nics -l`. 2 Remove the reported vmnic from the port groups by running `esxcfg-vswitch -U <affected vmnic#>` affected vSwitch. 3 Link available uplinks to the affected port groups by running `esxcfg-vswitch -L <available vmnic#>` affected vSwitch. Next, check the status of the vmnic in vSphere Client or the ESX service console: 1 In vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently unavailable. 2 From the service console, run `esxcfg-nics -l`. The output that appears is similar to the following example: Name PCI Driver Link Speed Duplex Description. ---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, shut down each

Alert Definition	Symptom	Recommendations
		ESX host port on the physical switch, run the "esxcfg-nics -l" command, and observe the affected vmnics. Verify that the vmnic identified in the alert is still connected to the switch and configured properly: 1 Make sure that the network cable is still connected to the switch and to the host. 2 Make sure that the switch is connected to the system, is still functioning properly, and was not inadvertently misconfigured. (See the switch documentation.) 3 Perform a network trace or observe activity LEDs to check for activity between the physical switch and the vmnic. 4 Check for network port settings on the physical switch. If the problem is caused by hardware, contact your hardware vendor for a hardware replacement.
A PCIe error occurred during system boot, but the error is recoverable.	A recoverable PCIe error occurred.	The PCIe error is recoverable, but the system behavior is dependent on how the error is handled by the OEM vendor's firmware. Contact the vendor for assistance.
A recoverable memory error has occurred on the host.	A recoverable memory error occurred.	Since recoverable memory errors are vendor-specific, contact the vendor for assistance.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
ESXi Host is violating vSphere 5.5 Hardening Guide.	■ Active directory authentication disabled OR ■ Non-compliant NTP service startup policy OR ■ SSH service is running OR ■ NTP service stopped OR ■ Non-compliant timeout value for automatically disabling local and remote shell access OR ■ vSphere Authentication Proxy not used for password protection when adding ESXi hosts to active directory OR ■ Persistent logging disabled OR ■ Bidirectional CHAP for iSCSI traffic disabled OR ■ Non-compliant firewall setting to restrict access to NTP client OR ■ NTP server for time synchronization not configured OR ■ Non-compliant ESXi Shell service startup policy OR ■ Non-compliant firewall setting to restrict access to SNMP server OR ■ ESXi Shell service is running OR ■ Non-compliant DCUI service startup policy OR ■ Dvfilter bind IP address configured OR ■ Non-compliant SSH service startup policy OR ■ DCUI service is running OR ■ Non-compliant idle time before an interactive shell is automatically logged out OR ■ Non-compliant DCUI access user list OR ■ Remote syslog is not enabled	Fix the vSphere 5.5 Hardening Guide Rules Violations according to the recommendations in the vSphere5 Hardening Guide

vRealize Automation 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. There are no alerts generated for network profiles.

The health and risk thresholds are as follows:

Health

- When 25%-50% of the child objects have health issues, the parent object will trigger an alert with a Warning health level.
- When 50%-75% of the child objects have health issues, the parent object will trigger an alert with an Immediate health level.
- When 75%-100% of the child objects have health issues, the parent object will trigger an alert with a Critical health level.

Risk

- When 25%-50% of the child objects have risk issues, the parent object will trigger an alert with a Warning risk level.
- When 50%-75% of the child objects have risk issues, the parent object will trigger an alert with an Immediate risk level.
- When 75%-100% of the child objects have risk issues, the parent object will trigger an alert with a Critical risk level.

vSAN Alerts Definitions

vRealize Operations Manager generates an alert if a problem occurs with the components in the storage area network that the vSAN adapter is monitoring.

Table 7-114. vSAN Alert Definitions

Alert	Alert Level	Affected Object Type	Description
Performance Service on vSAN cluster might be off or experience issues	Critical	vSAN Adapter Instance	Triggered when the vSphere Virtual SAN Performance Service is off or experiences issues for one of the vSAN-enabled cluster compute resources. Cleared by enabling Virtual SAN performance service in vSphere.
vSAN adapter instance failed to collect data from Virtual SAN Health Service. The health service might have issues.	Critical	vSAN Adapter Instance	Triggered when the vSAN adapter instance is unable to collect data from the vSphere Virtual SAN Health Service for one of the vSAN-enabled cluster compute resources. Check Virtual SAN health service settings in vSphere.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
vSAN Cluster disk space usage is approaching capacity	Warning	Cluster Compute Resource	Triggered when the disk usage in a vSAN cluster reaches 80% of capacity. Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster.
vSAN Cluster disk space capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the disk usage in a vSAN cluster reaches 95% of capacity. Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster.
vSAN Cluster flash read cache is approaching capacity	Warning	Cluster Compute Resource	Triggered when the Read Cache (RC) in the vSAN cluster reaches 80% of capacity. Cleared by adding flash storage to the read cache.
vSAN Cluster flash read cache capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the Read Cache (RC) in the vSAN cluster reaches 95% of capacity. Cleared by adding flash storage to the read cache.
vSAN Cluster virtual disk count is approaching capacity	Warning	Cluster Compute Resource	Triggered when the number of virtual disks per host in the vSAN cluster reaches 75% of capacity. Cleared by adding most hosts to the cluster.
vSAN Cluster virtual disk count capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the number of virtual disks per host in the vSAN cluster reaches 95% of capacity. Cleared by adding most hosts to the cluster.
vSAN Disk Group read cache hit rate is less than 90%	Warning	vSAN Disk Group	Triggered when the vSAN disk group read cache hit rate is less than 90%. Cleared by adding more cache to accommodate the workload.
vSAN Disk Group read cache hit rate is less than 90% and write buffer free space is less than 10%	Warning	vSAN Disk Group	Triggered when the vSAN disk group read cache hit rate is less than 90% and the vSAN disk group write buffer free space is less than 10%. Cleared by adding more flash capacity to the vSAN disk group.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
vSAN Host has no VMkernel NIC configured	Immediate	Host System	Triggered when vSAN host has no VMkernel NIC configured. Clears when the symptom disappears.
One or more physical disks on vSAN host is experiencing software state health issues	Critical	Host System	Triggered when one or more physical disks on vSAN host is experiencing software state health issues.
vSAN Performance Service is unable to communicate and retrieve statistics from host	Critical	Host System	Triggered when vSAN Performance Service is unable to communicate and retrieve statistics from host.
vSAN enabled hosts have inconsistent values for advanced configuration options	Critical	vSAN Cluster	Triggered when some advanced configuration settings have different values on different hosts in the vSAN cluster.
vSAN is disabled on the host	Critical	Host System	Triggered when vSAN is disabled on the host. vSAN objects health will be impacted If this host stores any vSAN data (For example: virtual machine objects) on its local disks.
One or more vSAN enabled hosts are not in the same IP subnet	Critical	Host System	Triggered when one or more vSAN enabled hosts are not in the same IP subnet.
Host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic	Critical	Host System	Triggered when host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic. **Note** Even if an ESXi host is part of the vSAN cluster, but is not contributing storage, it must still have a VMkernel NIC configured for vSAN traffic.
One or more hosts in the vSAN cluster have misconfigured multicast addresses	Critical	Host System	Triggered when one or more hosts in the vSAN cluster have misconfigured multicast addresses.
vSAN health service is not installed on the host	Critical	Host System	Triggered when vSAN health service is not installed on the host.
Host in a vSAN cluster has IP multicast connectivity issue	Critical	Host System	Triggered when host in a vSAN cluster has IP multicast connectivity issue. It means that multicast is most likely the root cause of a vSAN network partition.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
Host in a vSAN cluster has connectivity issues and vCenter Server does not know its state	Critical	Host System	Triggered when host in a vSAN cluster has connectivity issues and vCenter Server does not know its state.
vSAN disk group has incorrect deduplication and compression configuration	Critical	Host System	Triggered when vSAN disk group has incorrect deduplication and compression configuration.
vSAN witness host has an invalid preferred fault domain	Critical	Host System	Triggered when vSAN witness host has an invalid preferred fault domain.
vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster	Critical	Host System	Triggered when vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster.
Host has invalid unicast agent and impacting the health of vSAN Stretched Cluster	Critical	Host System	Triggered when host has invalid unicast agent and impacting the health of vSAN Stretched Cluster. An invalid unicast agent on the host can cause a communication malfunction with the witness host.
Storage I/O controller driver is not VMware certified	Critical	Host System	Triggered when stability and integrity of vSAN may be at risk as the storage I/O controller driver is not VMware certified.
Storage I/O controller is not compatible with the VMware Compatibility Guide	Critical	Host System	Triggered when vSAN environment may be at risk as the Storage I/O controller on the ESXi hosts that are participating in a vSAN cluster are not compatible with the VMware Compatibility Guide.
vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster	Critical	Host System	Triggered when vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster.
Unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster	Critical	Host System	Triggered when unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster.
The preferred fault domain is not set for the witness host in a vSAN Stretched cluster	Critical	Host System	Triggered when the preferred fault domain is not set for the witness host in a vSAN Stretched cluster and affecting the operations of vSAN Stretched cluster.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
vSAN Stretched cluster contains a witness host without a valid disk group	Critical	Host System	Triggered when vSAN Stretched cluster contains a witness host without a valid disk group. If the witness host does not have any disk claimed by vSAN then its fault domain is not available.
vSAN Stretched cluster has inconsistent configuration for Unicast agent	Critical	Host System	Triggered when vSAN Stretched cluster contains multiple unicast agents. This means multiple unicast agents were set on non-witness hosts.
vSAN Stretched cluster does not contain a valid witness host	Critical	Host System	Triggered when vSAN Stretched cluster does not contain a valid witness host. This affects the operations of vSAN Stretched cluster.
vSAN cluster has multiple network partitions	Critical	Host System	Triggered when vSAN cluster has multiple network partitions due to a network issue.
Witness host is a part of vSAN Stretched cluster	Critical	Host System	Triggered when witness host is a part of the vCenter cluster, which forms vSAN Stretched cluster.
Witness host resides in one of the data fault domains	Critical	Host System	Triggered when witness host resides in one of the data fault domains. This affects the operations of vSAN Stretched cluster.
vSAN cluster has unexpected hosts	Critical	Host System	Triggered when vSAN cluster has unexpected hosts.
vSAN is unable to retrieve the physical disk information from host	Critical	Host System	Triggered when vSAN is unable to retrieve the physical disk information from host. vSAN Health Service may not be working properly on this host.
vCenter Server has lost connection to a host that is part of a vSAN cluster.	Critical	Host System	Triggered when host that is part of a vSAN cluster is in disconnected state or not responding and vCenter Server does not know its state.
vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk	Critical	Host System	Triggered when vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks.	Critical	Host System	Triggered when vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.	Critical	Host System	Triggered when vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
vSAN is using a physical disk which has high congestion value	Critical	Host System	Triggered when vSAN is using a physical disk which has high congestion value. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
Disk format version of one or more vSAN disks is out of date	Critical	Host System	Triggered when disk format version of one or more vSAN disks is out of date and is not compatible with other vSAN disks. This can lead to problems in creating or powering on VMs, performance degradation, and EMM failures.
vSAN Cluster has multiple Stats DB objects which are creating conflicts and affecting vSAN Performance Service	Critical	vSAN Cluster	Triggered when vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service.
vSAN cluster has issues in electing stats master of vSAN Performance service	Critical	vSAN Cluster	Triggered when vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service.
CLOMD process on the host has issues and impacting the functionality of vSAN cluster	Critical	Host System	Triggered when CLOMD process on the host has issues and impacting the functionality of vSAN cluster.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
Number of vSAN components on a disk is reaching or has reached its limit	Critical	vSAN Cluster	Triggered when number of vSAN components on a disk is reaching or has reached its limit. This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations.
vSAN HCL DB auto updater is not working properly	Critical	vSAN Cluster	Triggered when vSAN HCL DB auto updater is not working properly. This means that vSAN cannot download and update its HCL DB automatically.
vSAN has encountered an issue while reading the metadata of a physical disk	Critical	Host System	Triggered when vSAN has encountered an issue while reading the metadata of a physical disk and cannot use this disk.
Overall health of the physical disks in a vSAN Cluster is impacted	Critical	Host System	Triggered when overall health of the physical disks in a vSAN Cluster is impacted. See the health status of each physical disk individually on all the hosts.
vSAN performance service statistics database object is reporting issues.	Critical	Host System	Triggered when vSAN performance service statistics database object is reporting issues.
Overall health of vSAN objects is reporting issues	Critical	vSAN Cluster	Triggered when overall health of vSAN objects is reporting issues.
Number of vSAN components on a host is reaching or has reached its limit	Critical	Host System	Triggered when number of vSAN components on a host is reaching or has reached its limit. This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations.
Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster	Critical	vSAN Cluster	Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster.
Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster	Critical	Host System	Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster.
Statistics collection of vSAN performance service is not working correctly	Critical	vSAN Cluster	Triggered when statistics collection of vSAN performance service is not working correctly. This means that statistics collection or writing statistics data to storage have failed for three consecutive intervals.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
After one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects	Critical	vSAN Cluster	Triggered when after one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects.
vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations	Critical	vSAN Cluster	Triggered when vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations.
Disk load variance between some vSAN disks exceeded the threshold value	Critical	vSAN Cluster	Triggered when disk load variance between some vSAN disks exceeded the threshold value. vSAN cannot perform the load balance properly.
Host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host	Critical	Host System	Triggered when host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host.
Storage I/O controller drivers is not supported with the current version of ESXi running on the host	Critical	Host System	Triggered when stability and integrity of vSAN may be at risk as the storage I/O controller driver is not supported with the current version of ESXi running on the host.
vSAN HCL DB is not up-to-date	Critical	vSAN Cluster	Triggered when vSAN HCL DB is not up-to-date.
vSAN cluster health checks are reporting issues	Critical	vSAN Cluster	Triggered when vSAN cluster health checks are reporting issues.
vSAN cluster Hardware Compatibility health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster Hardware Compatibility health checks are experiencing issues.
vSAN cluster Limits health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster Limits health checks are experiencing issues.
vSAN cluster network health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster network health checks are experiencing issues.
Performance Service on vSAN cluster might be off or experiencing issues	Critical	vSAN Cluster	Triggered when performance service on vSAN cluster might be off or experiencing issues. vROps will not be able to collect the data from vSAN cluster in this state.
vSAN Stretched cluster health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN Stretched cluster health checks are experiencing issues.

Table 7-114. vSAN Alert Definitions (continued)

Alert	Alert Level	Affected Object Type	Description
MTU check (ping with large packet size) has failed on vSAN host	Critical	Host System	Triggered when MTU check (ping with large packet size) has failed on vSAN environment due to some MTU misconfiguration in the vSAN network.
Basic (unicast) connectivity check (normal ping) has failed on vSAN host	Critical	Host System	Triggered when basic (unicast) connectivity check (normal ping) has failed on vSAN host due to network misconfiguration.
vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.	Critical	Host System	Triggered when vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.

Alerts in the vSphere Web Client

The vSphere Web Client displays the results of health tests for the following vSAN monitored groups:

- Network
- Physical disk
- Cluster
- Limits
- Data
- Hardware compatibility
- Performance Service
- Stretched Cluster (if enabled)

Each group contains several individual checks. If a check fails, the vSAN adapter issues a warning or error level alert. The alert indicates the host or cluster where the problem occurred and provides a recommendation to clear the alert. For a complete list of all vSAN health test alerts, see Knowledge Base article 2114803.

vSphere Distributed Port Group

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Port objects in your environment.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Critical

Alert Definition	Symptom	Recommendations
One or more ports are in link down state.	Symptoms include all of the following: ■ Port is connected ■ One or more ports are in a link down state	Verify that there is physical connectivity for the NICs on the host. Verify the admin status on the port.
One or more ports are experiencing network contention.	Port is experiencing dropped packets	Check if the packet drops are due to high CPU resource utilization or uplink bandwidth utilization. User vMotion to migrate the virtual machine that the port is attached to a different host.

Virtual Machine Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the virtual machine objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine is experiencing memory compression, ballooning or swapping due to memory limit.	■ Virtual machine memory limit is set AND ■ Virtual machine memory demand exceeds configured memory limit AND ■ [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/ immediate/critical level] AND ■ Recommended virtual machine memory size	Increase the memory limit for the virtual machine to match the recommended memory size. Alternatively, remove memory limit for the virtual machine.
Virtual machine has CPU contention caused by swap wait.	Virtual machine CPU swap wait is at warning/Immediate/Critical level.	1 Upgrade the host with more memory. 2 Use vSphere vMotion to migrate this virtual machine to a different host or cluster. 3 Set memory reservations for the virtual machine to prevent swapping.
Virtual machine has CPU contention caused by IO wait.	Virtual machine CPU I/O wait is at warning/immediate/critical level.	Increase the datastore I/O capacity for the connected data stores to reduce CPU I/O wait on the virtual machine.
Virtual machine has unexpected high CPU workload.	Symptoms include all of the following: ■ Virtual machine CPU demand at warning/immediate/critical level ■ Anomaly is starting to/ moderately/critically high	1 Check the guest applications to determine whether high CPU workload is an expected behavior. 2 Add more CPU capacity for this virtual machine.
Virtual machine has unexpected high memory workload.	Symptoms include all of the following: ■ Virtual machine memory workload is at Warning/ Immediate/Critical level ■ Anomaly is starting to/ moderately/critically high	1 Check the guest applications to determine whether high memory workload is an expected behavior. 2 Add more memory for this virtual machine.
Virtual machine has memory contention due to swap wait and high disk read latency.	Symptoms include all of the following: ■ Virtual machine CPU swap wait is at warning/immediate/critical level (5/10/15) ■ Virtual machine has read latency at warning level ■ Recommended virtual machine memory size	Add more memory for this virtual machine.

Alert Definition	Symptom	Recommendations
Virtual machine has memory contention due to memory compression, ballooning or swapping.	■ ! Virtual machine memory limit is set AND ■ Virtual machine has memory contention at warning/immediate/ critical level AN ■ [Virtual machine memory ballooning at warning/immediate/ critical level OR ■ Virtual machine memory is compressed OR ■ Virtual machine is using swap]	1 Add memory reservations to this virtual machine to prevent ballooning and swapping. 2 Use vSphere vMotion to migrate this virtual machine to a different host or cluster.
Virtual machine has unexpected high disk I/O workload.	Symptoms include all of the following: ■ Virtual machine disk I/O workload at Warning/Immediate/Critical level (80/90/95) ■ Virtual machine disk I/O workload above DT	1 Check the applications running on the virtual machine to determine whether high disk I/O workload is an expected behavior. 2 Use vSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS.
Virtual machine has disk I/O read latency problem.	Symptoms include all of the following: ■ Virtual machine disk read latency at Warning /Immediate/Critical level ■ Virtual machine disk read latency above DT ■ Virtual machine has low co-stop ■ Virtual machine has low CPU swap wait	1 Check whether you have enabled Storage IO control on the datastores connected to the virtual machine. 2 Increase IOPS for the datastores connected to the virtual machine. 3 UsevSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS.
Virtual machine has disk I/O write latency problem.	Symptoms include all of the following: ■ Virtual machine disk write latency at Warning/ Immediate/Critical level ■ Virtual machine disk write latency above DT ■ Virtual machine has low CPU swap wait (< 3 ms)	1 Check whether you have enabled Storage IO Control on the data stores connected to the datastore. 2 Increase IOPS for the data stores connected to the virtual machine. 3 If the virtual machine has multiple snapshots, delete the older snapshots. 4 Use vSphere Storage vMotion to migrate some virtual machines to a different datastore.
Virtual machine has disk I/O latency problem caused by snapshots.	Symptoms include all of the following: ■ Virtual machine CPU I/O wait is at warning/immediate/critical level ■ Virtual machine has at least one snapshot ■ All child datastores have [! Disk command latency at warning level]	1 If the virtual machine has multiple snapshots, delete the older snapshots. 2 Reduce the number of snapshots by consolidating the snapshots into one snapshot. In vSphere Client, select the VM, right-click, select **Snapshot**, and then **Consolidate**.

Alert Definition	Symptom	Recommendations
Virtual machine is consuming disk space in a rapid and unexpected manner.	Symptoms include all of the following: ■ Guest file system overall disk space usage reaching warning/ immediate/critical limit (80, 90, 95) ■ Virtual machine disk space time remaining high (> 60 days) ■ Guest file system space usage above DT ■ Guest partition disk space usage	1 Check the application and verify that it is behaving correctly. 2 Add a new hard disk to the virtual machine and configure the guest file system partition to use the disk.
One or more guest file systems is out of disk space.	One or more guest file systems out of disk space (Fault symptom).	Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk.
Not enough resources for vSphere HA to start the virtual machine.	Not enough resources forvSphere HA to start VM (Fault symptom).	1 If virtual machine CPU reservation is set, decrease the CPU reservation configuration. 2 If virtual machine memory reservation is set, decrease the memory reservation configuration. 3 Add more hosts to cluster. 4 Bring any failed hosts online or resolve a network partition, if one exists. 5 If DRS is in manual mode, look for pending recommendations and approve the recommendations so that vSphere HA failover can proceed.
The Fault tolerance state of the virtual machine has changed to "Disabled" state.	VM fault tolerance state changed to disabled (Fault symptom).	Enable the secondary virtual machine indicated in the alert.
vSphere HA failed to restart a network isolated virtual machine.	vSphere HA failed to restart a network isolated virtual machine (Fault symptom).	Manually power on the virtual machine.
The fault tolerance state of the virtual machine has changed to "Needs Secondary" state.	VM Fault Tolerance state changed to needs secondary (Fault symptom).	Keep HA enabled when Fault tolerance (FT) is required to protect virtual machines.

Alert Definition	Symptom	Recommendations
vSphere HA cannot perform a failover operation for the virtual machine	vSphere HA virtual machine failover unsuccessful (Fault symptom)	1 If the error information reports that a file is locked, the virtual machine might be powered on a host that the vSphere HAmaster agent can no longer monitor by using the management network or heartbeat datastores. 2 The virtual machine might have been powered on by a user on a host outside of the cluster. If any hosts are declared offline, determine whether a networking or storage problem caused the situation. 3 If the error information reports that the virtual machine is in an invalid state, an in-progress operation might be preventing access to the virtual machine files. Determine whether any operations are in progress, such as a clone operation that is taking a long time to complete. 4 You can also try to power on the virtual machine and investigate any returned errors.
Virtual machine is experiencing memory compression, ballooning or swapping due to memory limit.	■ Virtual machine memory limit is set ■ Virtual machine memory demand exceeds configured memory limit ■ [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/ immediate/critical level] ■ Recommended virtual machine memory size	Increase the memory limit for the virtual machine to match the recommended memory size. Alternatively, remove memory limit for the virtual machine.

Efficiency/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Efficiency

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine has large disk snapshots.	Symptoms include all of the following: ■ Virtual machine has large disk snapshots ■ Reclaimable snapshot waste ■ Datastore space usage reaching warning/immediate/critical limit	If the virtual machine has multiple snapshots, delete the older snapshots.

Efficiency/Warning

These alert definitions have the following impact and criticality information.

Impact

Efficiency

Criticality

Warning

Alert Definition	Symptom	Recommendations
Virtual machine is idle.	Symptoms include all of the following: ■ Virtual machine is idle ■ Virtual machine high ready time on each vCPU ■ ! Virtual machine is powered off	Power off this virtual machine to allow for other virtual machines to use CPU and memory that this virtual machine is wasting.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine has CPU contention caused by co-stop.	Symptoms include all of the following: ■ Virtual machine CPU co-stop at warning/immediate/critical level ■ ! Virtual machine is powered off ■ Number of vCPUs to remove from virtual machine	Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the symptom.
Virtual machine has chronic high CPU workload leading to CPU stress.	Symptoms include all of the following: ■ Virtual machine CPU stress is at warning/immediate/critical level ■ Recommended number of vCPUs to add	Add more CPU capacity for this virtual machine.
Virtual machine has high CPU co-stop due to snapshots.	Symptoms include all of the following: ■ Virtual machine CPU co-stop is at warning/immediate/critical level ■ Virtual machine has at least one snapshot	To reduce the high co-stop (%CSTP) values and increase virtual machine performance, consolidate any snapshots into the main virtual disk. In the vSphere Client, select the VM, right click, and select **Snapshot**, and then **Consolidate**. After consolidation, the %CSTP value is reduced or eliminated and VM performance is improved. If performance is not improved enough, continue researching other potential VM performance issues. See VMware KB: http://kb.vmware.com/kb/2000058
Virtual machine has chronic high memory workload leading to memory stress.	Symptoms include all of the following: ■ Virtual machine memory stress at warning/immediate/critical level ■ Recommended virtual machine memory size > 0	Add more memory for the VM.
Virtual machine is projected to run out of disk space.	Symptoms include all of the following: ■ Virtual machine disk space time remaining low (<= 60 days) ■ ! Guest file system space usage above DT ■ ! Guest file system overall disk space usage reaching warning limit (85%) ■ Guest partition disk space usage	1 Check the application configuration to determine whether the virtual machine disk capacity will be sufficient. 2 Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk.

Alert Definition	Symptom	Recommendations
Virtual machine is running out of disk space.	Symptoms include all of the following: ■ Guest file system overall disk space usage reaching warning/ immediate/critical limit (80, 90, 95) ■ Virtual machine disk space time remaining low (<= 60 days) ■ ! Guest file system space usage above DT ■ Guest partition disk space usage	1 Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk. 2 Reclaim disk space using in-guest disk cleanup mechanisms.
Virtual machine is violating vSphere 5.5 hardening guide.	■ Unrestricted VM-to-VM communication through VMCI OR ■ VMsafe CPU/Memory APIs-port number configured OR ■ Dvfilter network API enabled OR ■ Non-compliant max VMX file size OR ■ Non-compliant max VM log file size OR ■ Allow unauthorized modification of device settings OR ■ Allow unauthorized connect and disconnect of devices OR ■ Tools auto install not disabled OR ■ Non-compliant max number of remote console connections OR ■ Allow VM to obtain detailed information about the physical host OR ■ Non-compliant max VM log file count OR ■ Feature not exposed in vSphere: MemsFss is not disabled OR ■ VMsafe CPU/memory API enabled OR ■ Parallel port connected OR ■ Console drag and drop operation not disabled OR ■ Console copy operation not disabled OR ■ Serial port connected OR ■ Feature not exposed in vSPhere: AutoLogon is not disabled OR ■ Use independent non persistent disk OR ■ Feature not exposed in vSphere: UnityPush is not disabled OR	Fix the vSphere 5.5 hardening guide rule violations according to the recommendations in the vSphere Hardening Guide (XLSX).

Alert Definition	Symptom	Recommendations
	■ Shrink virtual disk not disabled - diskShrink OR ■ Feature not exposed in vSphere: GetCreds is not disabled OR ■ CD-ROM connected OR ■ Feature not exposed in vSPhere: HGFSServerSet is not disabled OR ■ Console paste operation not disabled OR ■ Feature not exposed in vSPhere: BIOSBBS is not disabled OR ■ Shrink virtual disk not disabled - diskWiper OR ■ USB controller connected OR ■ Feature not exposed in vSPhere: Monitor Control is not disabled OR ■ Floppy drive connected OR ■ Feature not exposed in vSphere: LaunchMenu is not disabled OR ■ Versionget is not disabled OR ■ Feature not exposed in vSphere: Toporequest is not disabled OR ■ Feature not exposed in vSphere: Unity-interlock not disabled OR ■ VM logging is not disabled OR ■ Feature not exposed in vSphere: Unity is not disabled OR ■ Feature not exposed in vSphere: Trashfolderstate is not disabled OR ■ VGA only mode is not enabled OR ■ Feature not exposed in vSphere: Trayicon is not disabled OR ■ Feature not exposed in vSphere: Unity-Taskbar is not disabled OR ■ Feature not exposed in vSphere: Versionset is not disabled OR ■ VM console access via VNC protocol is not disabled OR ■ Feature not exposed in vSphere: Protocolhandler is not disabled OR ■ VIX message is not disabled OR ■ Feature not exposed in vSphere: Shellaction is not disabled OR ■ 3D features is not disabled OR	

Alert Definition	Symptom	Recommendations
	■ Feature not exposed in vSphere: Unity-Windowcontents is not disabled OR ■ Feature not exposed in vSphere: Unity-Unityactive is not disabled	

Risk/Warning

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Warning

Alert Definition	Symptom	Recommendations
Virtual machine is demanding more CPU than the configured limit.	Symptoms include all of the following: ■ Virtual machine CPU limit is set ■ Virtual machine CPU demand exceeds configured limit ■ ! Virtual machine's CPU demand exceeds its provisioned capacity	Increase of remove CPU limits on the VM.

vSphere Distributed Switch Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Switch objects in your environment.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Critical

Alert Definition	Symptom	Recommendations
Network traffic is blocked for one or more ports.	Network traffic is blocked for one or more ports.	Check the security policy on the port groups as well as any ACL rule configuration.

Health/Warning

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Warning

Alert Definition	Symptom	Recommendations
Distributed Switch configuration is out of sync.	Distributed Switch configuration is out of sync with the vCenter Server.	Change the distributed switch configuration to match the host. Identify the distributed switch properties that are out of sync. If these properties were changed locally on the host in order to maintain connectivity, update the distributed switch configuration in the vCenter Server. Otherwise, re-apply the the vCenter Server configuration to this host.
One or more VLANs are unsupported by the physical switch.	One or more VLANs are unsupported by the physical switch.	Ensure the VLAN configuration on the physical switch and the distributed port groups are consistent.
Teaming configuration does not match the physical switch.	Teaming configuration does not match the physical switch.	Ensure the teaming configuration on the physical switch and the distributed switch are consistent.
The MTU on the Distributed Switch is not allowed by one or more VLANs on the host.	The MTU on the Distributed Switch is not allowed by one or more VLANs on the host.	Ensure the MTU configuration on the physical switch and the distributed switch are consistent.
There is an MTU mismatch between the host and a physical switch.	There is an MTU mismatch between the host and a physical switch.	Adjust the MTU configuration on the host to match the physical switch. Change the MTU configuration on the physical switch.

Risk/Warning

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Warning

Alert Definition	Symptom	Recommendations
The distributed switch configuration is incorrect.	Host without redundant physical connectivity to the distributed switch.	Verify that at least two NICs on each host is connected to the distributed switch.

vCenter Server Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vCenter Server objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
A problem occurred with a vCenter Server component.	The vCenter Server health changed (fault symptom).	The actions to take to resolve the problems depend on the specific problem that caused the fault. Review the issue details, and check the documentation.
Duplicate object name found in thevCenter Server.	Duplicate object name found in the vCenter Server.	Ensure the virtual machines names are unique before enabling the Name-Based Identification feature.
The vCenter Server Storage data collection failed.	The vCenter Server storage data collection failed.	Ensure vCenter Management Webservice is started and Storage Management Service is functioning.

Datastore Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
Datastore has unexpected high Disk I/O workload.	Symptoms include all of the following: ■ Datastore disk I/O workload at warning/immediate/critical level ■ Datastore disk I/O workload above DT	1 Check the applications running on the virtual machines placed on the datastore to determine whether high disk I/O workload is expected behavior. 2 Increase IOPS for the datastore.
Datastore is consuming disk space in a rapid and unexpected manner.	Symptoms include all of the following: ■ Datastore space usage reaching warning/immediate/critical level ■ Datastore space growth above DT ■ Datastore time remaining high	1 Check if there is an unexpected provisioning of virtual machines on this datastore. 2 Use vSphere Storage vMotion to migrate some virtual machines to a different datastore. 3 Add more capacity to the datastore.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Critical

Alert Definition	Symptom	Recommendations
A storage device for a datastore has been detected to be off.	Storage device has been turned off administratively (fault symptom).	Ask the administrator about the device state. The fault will be resolved and the alert canceled if the device is turned on. If SCSI devices are detached or permanently removed, you must manually cancel the alert.
Datastore has lost connectivity to a storage device.	Host(s) lost connectivity to storage device(s) (fault symptom).	The storage device path, for example, `vmhba35:C1:T0:L7`, contains several potential failure points: Path Element | Failure Point --- vmhba35 | HBA (Host Bus Adapter) C1 | Channel T0 | Target (storage processor port) L7 | LUN (Logical Unit Number or Disk Unit). To determine the cause of the failure or to eliminate possible problems: Identify the available storage paths to the reported storage device by running `esxcfg-mpath -l`. For more information, see https://kb.vmware.com/s/article/1003973?lang=en_US#q=1003973. Check that a rescan does not restore visibility to the targets. For information on rescanning the storage device by using the command-line interface and the vSphere Client, see https://kb.vmware.com/s/article/1003988?lang=en_US#q=1003988. Determine whether the connectivity issue is with the iSCSI storage or the fiber storage. Troubleshoot the connectivity to the iSCSI storage by using the software initiator: 1 Check whether a ping to the storage array fails from ESX. For more information, see https://kb.vmware.com/s/article/1003486?lang=en_US#q=1003486 2 Check whether a vmkping to each network portal of the storage array fails. For more information, see https://kb.vmware.com/s/article/1003728?lang=en_US#q=1003728. 3 Check that the initiator is registered on the array. For more information, contact your storage vendor.

Alert Definition	Symptom	Recommendations
		4 Check that the following physical hardware is functioning correctly: Ethernet switch, Ethernet cables between the switch and the ESX host, and Ethernet cables between the switch and the storage array. To troubleshoot the connectivity to the fiber-attached storage, check the fiber switch. The fiber switch zoning configuration permits the ESX host to see the storage array. If you require assistance, contact your switch vendor. The fiber switch propagates RSCN messages to the ESX hosts. For more information about configuring the fiber switch, see https://kb.vmware.com/s/article/1002301?lang=en_US#q=1002301. Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself. You must rescan after making changes to make sure that the targets are detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or change, you must cancel the fault alert as a workaround. The alert will then be canceled automatically.

Health/Immediate

These alert definitions have the following impact and criticality information.

Impact

Health

Criticality

Immediate

Alert Definition	Symptom	Recommendations
Datastore has one or more hosts that have lost redundant paths to a storage device.	Host(s) lost redundancy to storage device(s) (fault symptom).	The storage device path, for example, vmhba35:C1:T0:L7, contains several potential failure points: Path Element | Failure Point --- vmhba35 | HBA (Host Bus Adapter) C1 | Channel T0 | Target (storage processor port) L7 | LUN (Logical Unit Number or Disk Unit). Use the following guidance to determine the cause of the failure or to eliminate possible problems. Identify the available storage paths to the reported storage device by running `esxcfg-mpath -l`. For more information, see https://kb.vmware.com/s/article/1003973?lang=en_US#q=1003973. Check that a rescan does not restore visibility to the targets. For information on rescanning the storage device by using the command-line interface and the vSphere Client, see https://kb.vmware.com/s/article/1003988?lang=en_US#q=1003988. Determine whether the connectivity issue is with the iSCSI storage or the fiber storage. Troubleshoot the connectivity to the iSCSI storage by using the software initiator: 1 Check whether a ping to the storage array fails from ESX. For more information, see https://kb.vmware.com/s/article/1003486?lang=en_US#q=1003486. 2 Check whether a vmkping to each network portal of the storage array fails. For more information, see https://kb.vmware.com/s/article/1003728?lang=en_US#q=1003728. 3 Check that the initiator is registered on the array. For more information, contact your storage vendor. 4 Check that the following physical hardware is functioning correctly: Ethernet switch, Ethernet cables between the switch and the ESX host, and Ethernet cables between the switch and the storage array.

Alert Definition	Symptom	Recommendations
		To troubleshoot the connectivity to the fiber-attached storage, check the fiber switch. The fiber switch zoning configuration permits the ESX host to see the storage array. If you require assistance, contact your switch vendor. The fiber switch propagates RSCN messages to the ESX hosts. For more information about configuring the fiber switch, see https://kb.vmware.com/s/article/1002301?lang=en_US#q=1002301. Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself. You must rescan after making changes to make sure that the targets are detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or change, you must cancel the fault alert as a workaround. The alert will be canceled automatically after that.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptom	Recommendations
Datastore is running out of disk space.	Symptoms include all of the following: ■ Datastore space usage reaching warning/immediate/critical level ■ ! Datastore space growth above DT ■ Datastore space time remaining is low	1 Add more capacity to the datastore. 2 Use vSphere vMotion to migrate some virtual machines to a different datastore. 3 Delete unused snapshots of virtual machines from datastore. 4 Delete any unused templates on the datastore.
Datastore is projected to run out of disk space.	Symptoms include all of the following: ■ ! Datastore space usage reaching warning level ■ ! Datastore space growth above DT ■ Datastore space time remaining is low	1 Check if datastore usage is a planned growth and expand the storage if necessary. 2 Use vSphere vMotion to migrate some virtual machines to a different datastore.

Data Center Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Data Center objects in your environment.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information:

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Data center has unbalanced CPU "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ DC is unbalanced on CPU "demand" workload ■ DC has significant CPU "demand" workload difference ■ At least one cluster in DC has high CPU "demand" workload	Rebalance the container to spread the workload more evenly.
Data center has unbalanced memory "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully enabled ■ DC is unbalanced on memory "demand" workload difference ■ At least one cluster in DC has high memory "demand" workload	Rebalance the container to spread the workload more evenly.
Data center has unbalanced memory "consumed" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ DC is unbalanced on memory "consumed" workload ■ DC has significant memory "consumed" workload difference ■ At least one cluster in DC has high memory "consumed" workload	Rebalance the container to spread the workload more evenly.

Custom Data Center Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Custom Data Center objects in your environment.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Custom data center has unbalanced CPU "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on CPU "demand" workload ■ CDC has significant CPU "demand" workload difference ■ At least one cluster in CDC has high CPU "demand" workload	Rebalance the container to spread the workload more evenly.
Custom data center has unbalanced memory "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on memory "demand" workload ■ CDC has significant memory "demand" workload difference ■ At least one cluster in CDC has high memory"demand" workload	Rebalance the container to spread the workload more evenly.
Custom Datacenter has unbalanced memory "consumed" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on memory "consumed" workload ■ CDC has significant memory "consumed" workload difference ■ At least one cluster in CDC has high memory"consumed" workload	Rebalance the container to spread the workload more evenly.

Property Definitions in vRealize Operations Manager

Properties are attributes of objects in the vRealize Operations Manager environment. You use properties in symptom definitions. You can also use properties in dashboards, views, and reports.

vRealize Operations Manager uses adapters to collect properties for target objects in your environment. Property definitions for all objects connected through the vCenter adapter are provided. The properties collected depend on the objects in your environment.

You can add symptoms based on properties to an alert definition so that you are notified if a change occurs to properties on your monitored objects. For example, disk space is a hardware property of a virtual machine. You can use disk space to define a symptom that warns you when the value falls below a certain numeric value. See Defining Symptoms for Alerts.

vRealize Operations Manager generates Object Type Classification and Subclassification properties for every object. You can use object type classification properties to identify whether an object is an adapter instance, custom group, application, tier, or a general object with property values *ADAPTER_INSTANCE*, *GROUP*, *BUSINESS_SERVICE*, *TIER*, or *GENERAL*, respectively.

Properties for vCenter Server Components

The VMware vSphere solution is installed with vRealize Operations Managerand includes the vCenter adapter. vRealize Operations Manager uses the vCenter adapter to collect properties for objects in the vCenter Server system.

vCenter Server components are listed in the `describe.xml` file for the vCenter adapter. The following example shows the runtime property `memoryCap` or Memory Capacity for the virtual machine in the `describe.xml`.

```
<ResourceGroup instanced="false" key="runtime" nameKey="5300" validation="">
   <ResourceAttribute key="memoryCap" nameKey="1780" dashboardOrder="200" dataType="float"
                      defaultMonitored="true" isDiscrete="false" isRate="false" maxVal=""
                      minVal="" isProperty="true" unit="kb"/>
</ResourceGroup>
```

The `ResourceAttribute` element includes the name of the property that appears in the UI and is documented as a Property Key. `isProperty = "true"` indicates that `ResourceAttribute` is a property.

vCenter Server Properties

vRealize Operations Manager collects summary and event properties for vCenter Server system objects.

Table 7-115. Summary Properties Collected for vCenter ServerSystem Objects

Property Key	Property Name	Description
summary|version	Version	Version
summary|vcuuid	VirtualCenter ID	Virtual Center ID
summary|vcfullname	Product Name	Product Name

Table 7-116. Event Properties Collected for vCenter ServerSystem Objects

Property Key	Property Name	Description
event|time	Last VC Event Time	Last Virtual Center Event Time
event|key	Last VC Event ID	Last Virtual Center Event ID

Table 7-117. Custom Field Manager Property Collected for vCenter ServerSystem Objects

Property Key	Property Name	Description
CustomFieldManager|CustomFieldDef	Custom Field Def	Custom Field Def for VCenter Tagging information at Adapter level.

Virtual Machine Properties

vRealize Operations Manager collects configuration, runtime, CPU, memory, network I/O, summary, guest file system, and properties about datastore use for virtual machine objects.

Table 7-118. vRealize Automation Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
vRealize Automation|Blueprint Name	Blueprint Name	Virtual machines deployed byvRealize Automation to be excluded from workload placements.

Table 7-119. Properties Collected for Virtual Machine Objects to Support VIN Adapter Localization

Property Key	Property Name	Description
RunsOnApplicationComponents	Application components running on the Virtual Machine	Application components running on the Virtual Machine
DependsOnApplicationComponents	Application components the Virtual Machine depends on	Application components running on other machines that this Virtual Machine depends on.

Table 7-120. Configuration Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
config|name	Name	Name
config|guestFullName	Guest Fullname	Guest OS full name configured by the user.
config|hardware|numCpu	Number of virtual CPUs	Number of virtual CPUs
config|hardware|memoryKB	Memory	Memory
config|hardware|thinEnabled	Thin Provisioned Disk	Indicates whether thin provisioning is enabled
config|hardware|diskSpace	Disk Space	Disk Space
config|cpuAllocation|reservation	Reservation	CPU reservation
config|cpuAllocation|limit	Limit	CPU limit
config|cpuAllocation|shares|shares	Shares	CPU shares
config|memoryAllocation|reservation	Reservation	CPU reservation
config|memoryAllocation|limit	Limit	Limit
config|memoryAllocation|shares|shares	Shares	Memory shares
config|extraConfig|mem_hotadd	Memory Hot Add	Memory Hot Add Configuration
config|extraConfig|vcpu_hotadd	VCPU Hot Add	VCPU Hot Add Configuration
config|extraConfig|vcpu_hotremove	VCPU Hot Remove	VCPU Hot Remove Configuration
config|security|disable_autoinstall	Disable tools auto install (isolation.tools.autoInstall.disable)	Disable tools auto install (isolation.tools.autoInstall.disable)

Table 7-120. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
config|security|disable_console_copy	Disable console copy operations (isolation.tools.copy.disable)	Disable console copy operations (isolation.tools.copy.disable)
config|security|disable_console_dnd	Disable console drag and drop operations (isolation.tools.dnd.disable)	Disable console drag and drop operations (isolation.tools.dnd.disable)
config|security|enable_console_gui_options	Enable console GUI operations (isolation.tools.setGUIOptions.enable)	Enable console GUI operations (isolation.tools.setGUIOptions.enable)
config|security|disable_console_paste	Disable console paste operations (isolation.tools.paste.disable)	Disable console paste operations (isolation.tools.paste.disable)
config|security|disable_disk_shrinking_shrink	Disable virtual disk shrink (isolation.tools.diskShrink.disable)	Disable virtual disk shrink (isolation.tools.diskShrink.disable)
config|security|disable_disk_shrinking_wiper	Disable virtual disk wiper (isolation.tools.diskWiper.disable)	Disable virtual disk wiper (isolation.tools.diskWiper.disable)
config|security|disable_hgfs	Disable HGFS file transfers (isolation.tools.hgfsServerSet.disable)	Disable HGFS file transfers (isolation.tools.hgfsServerSet.disable)
config|security|disable_independent_nonpersistent	Avoid using independent nonpersistent disks (scsiX:Y.mode)	Avoid using independent nonpersistent disks (scsiX:Y.mode)
config|security|enable_intervm_vmci	Enable VM-to-VM communication through VMCI (vmci0.unrestricted)	Enable VM-to-VM communication through VMCI (vmci0.unrestricted)
config|security|enable_logging	Enable VM logging (logging)	Enable VM logging (logging)
config|security|disable_monitor_control	Disable VM Monitor Control (isolation.monitor.control.disable)	Disable VM Monitor Control (isolation.monitor.control.disable)
config|security|enable_non_essential_3D_features	Enable 3D features on Server and desktop virtual machines (mks.enable3d)	Enable 3D features on Server and desktop virtual machines (mks.enable3d)
config|security|disable_unexposed_features_autologon	Disable unexposed features - autologon (isolation.tools.ghi.autologon.disable)	Disable unexposed features - autologon (isolation.tools.ghi.autologon.disable)
config|security|disable_unexposed_features_biosbbs	Disable unexposed features - biosbbs (isolation.bios.bbs.disable)	Disable unexposed features - biosbbs (isolation.bios.bbs.disable)
config|security|disable_unexposed_features_getcreds	Disable unexposed features - getcreds (isolation.tools.getCreds.disable)	Disable unexposed features - getcreds (isolation.tools.getCreds.disable)
config|security|disable_unexposed_features_launchmenu	Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change)	Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change)

Table 7-120. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
config|security|disable_unexposed_features_memsfss	Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable)	Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable)
config|security|disable_unexposed_features_protocolhandler	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)
config|security|disable_unexposed_features_shellaction	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)
config|security|disable_unexposed_features_toporequest	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)
config|security|disable_unexposed_features_trashfolderstate	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)
config|security|disable_unexposed_features_trayicon	Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable)	Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable)
config|security|disable_unexposed_features_unity	Disable unexposed features - unity (isolation.tools.unity.disable)	Disable unexposed features - unity (isolation.tools.unity.disable)
config|security|disable_unexposed_features_unity_interlock	Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable)	Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable)
config|security|disable_unexposed_features_unity_taskbar	Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable)	Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable)
config|security|disable_unexposed_features_unity_unityactive	Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable)	Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable)
config|security|disable_unexposed_features_unity_windowcontents	Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable)	Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable)
config|security|disable_unexposed_features_unitypush	Disable unexposed features - unitypush (isolation.tools.unity.push.update.disable)	Disable unexposed features - unitypush (isolation.tools.unity.push.update.disable)

Table 7-120. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
config|security|disable_unexposed_features_versionget	Disable unexposed features - versionget (isolation.tools.vmxDnDVersionGet.disable)	Disable unexposed features - versionget (isolation.tools.vmxDnDVersionGet.disable)
config|security|disable_unexposed_features_versionset	Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable)	Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable)
config|security|disable_vix_messages	Disable VIX messages from the VM (isolation.tools.vixMessage.disable)	Disable VIX messages from the VM (isolation.tools.vixMessage.disable)
config|security|enable_vga_only_mode	Disable all but VGA mode on virtual machines (svga.vgaOnly)	Disable all but VGA mode on virtual machines (svga.vgaOnly)
config|security|limit_console_connection	Limit number of console connections (RemoteDisplay.maxConnection)	Limit number of console connections (RemoteDisplay.maxConnection)
config|security|limit_log_number	Limit number of log files (log.keepOld)	Limit number of log files (log.keepOld)
config|security|limit_log_size	Limit log file size (log.rotateSize)	Limit log file size (log.rotateSize)
config|security|limit_setinfo_size	Limit VMX file size (tools.setInfo.sizeLimit)	Limit VMX file size (tools.setInfo.sizeLimit)
config|security|enable_console_VNC	Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled)	Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled)
config|security|disable_device_interaction_connect	Disable unauthorized removal, connection of devices (isolation.device.connectable.disable)	Disable unauthorized removal, connection of devices (isolation.device.connectable.disable)
config|security|disable_device_interaction_edit	Disable unauthorized modification of devices (isolation.device.edit.disable)	Disable unauthorized modification of devices (isolation.device.edit.disable)
config|security|enable_host_info	Enable send host information to guests (tools.guestlib.enableHostInfo)	Enable send host information to guests (tools.guestlib.enableHostInfo)
config|security|network_filter_enable	Enable dvfilter network APIs (ethernetX.filterY.name)	Enable dvfilter network APIs (ethernetX.filterY.name)
config|security|vmsafe_cpumem_agentaddress	VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress)	VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress)
config|security|vmsafe_cpumem_agentport	VMsafe CPU/memory APIs - port number (vmsafe.agentPort)	VMsafe CPU/memory APIs - port number (vmsafe.agentPort)
config|security|vmsafe_cpumem_enable	Enable VMsafe CPU/memory APIs (vmsafe.enable)	Enable VMsafe CPU/memory APIs (vmsafe.enable)
config|security|disconnect_devices_floppy	Disconnect floppy drive	Disconnect floppy drive

Table 7-120. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
config|security|disconnect_devices_cd	Disconnect CD-ROM	Disconnect CD-ROM
config|security|disconnect_devices_usb	Disconnect USB controller	Disconnect USB controller
config|security|disconnect_devices_parallel	Disconnect parallel port	Disconnect parallel port
config|security|disconnect_devices_serial	Disconnect serial port	Disconnect serial port

Note Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

For more information on the *vSphere Hardening Guide* alerts, see Customize a Policy to Enable the vSphere Hardening Guide Alerts.

Table 7-121. Runtime Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
runtime|memoryCap	Memory Capacity	Memory Capacity

Table 7-122. CPU Usage Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
cpu|limit	CPU limit	CPU limit
cpu|reservation	CPU reservation	CPU reservation
cpu|speed	CPU	CPU Speed
cpu|cpuModel	CPU Model	CPU Model

Table 7-123. Memory Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
mem|host_reservation	VM Reservation	Mem Machine Reservation
mem|host_limit	VM Limit	Mem Machine Limit

Table 7-124. Network Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
net|mac_address	Mac Address	Mac Address
net|ip_address	IP Address	IP Address
net|subnet_mask	Subnet Mask	Subnet Mask
net|default_gateway	Default Gateway	Default Gateway
net|nvp_vm_uuid	NVP VM UUID	NVP VM UUID

Table 7-125. Summary Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name
summary|parentCluster	Parent Cluster	Parent Cluster
summary|parentHost	Parent Host	Parent Host
summary|parentDatacenter	Parent Datacenter	Parent Datacenter
summary|parentVcenter	Parent Vcenter	Parent Vcenter
summary|guest|fullName	Guest OS Full Name	Guest OS Full Name as identified by VMware tools
summary|guest|ipAddress	Guest OS IP Address	Guest OS IP Address
summary|guest|toolsRunningStatus	Tools Running Status	Guest Tools Running Status
summary|guest|toolsVersionStatus2	Tools Version Status	Guest Tools Version Status 2
summary|guest| vrealize_operations_agent_id	vRealize Operations Agent ID	An ID to identify a VM in Agent Adapter's world
summary|guest| vrealize_operations_euc_agent_id	vRealize Operations Euc Agent ID	An ID to identify a VM in Agent Adapter's world
summary|config|numEthernetCards	Number of NICs	Number of NICs
summary|config|isTemplate	VM Template	Indicates whether it is a VM Template
summary|runtime|powerState	Power State	Power State
summary|runtime|connectionState	Connection State	Connection State

Table 7-126. Datastore Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
datastore| maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore| maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Table 7-127. Guest File System Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
guestfilesystem|capacity_property	Guest File System Capacity Property	Total capacity of guest file system as a property, reported for each file system.
guestfilesystem| capacity_property_total	Total Guest File System Capacity Property	Overall total capacity of guest file system as a property, reported across all file systems.

Host System Properties

vRealize Operations Manager collects configuration, hardware, runtime, CPU, network I/O, summary, and properties about datastore use for host system objects.

Table 7-128. Configuration Properties Collected for Host System Objects

Property Key	Property Name	Description
config|name	Name	Name
config|diskSpace	Disk Space	Disk Space
config|network|nnic	Number of NICs	Number of NICs
config|network|linkspeed	Average Physical NIC Speed	Average Physical NIC Speed
config|network|dnsserver	DNS Server	List of DNS Servers
config|product|productLineId	Product Line ID	Product Line ID
config|product|apiVersion	API Version	API Version
config|storageDevice|plugStoreTopology|numberofPath	Total number of Path	Total number of storage paths
config|storageDevice|multipathInfo|numberofActivePath	Total number of Active Path	Total number of active storage paths
config|storageDevice|multipathInfo|multipathPolicy	Multipath Policy	Multipath Policy
config|hyperThread|available	Available	Indicates whether hyperthreading is supported by the server
config|hyperThread|active	Active	Indicates whether hyperthreading is active
config|ntp|server	NTP Servers	NTP Servers
config|security|ntpServer	NTP server	NTP server
config|security|enable_ad_auth	Enable active directory authentication	Enable active directory authentication
config|security|enable_chap_auth	Enable mutual chap authentication	Enable mutual chap authentication
config|security|enable_auth_proxy	Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate)	Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate)
config|security|syslog_host	Remote log host (Syslog.global.logHost)	Remote log host (Syslog.global.logHost)
config|security|dcui_access	Users who can override lock down mode and access the DCUI (DCUI.Access)	Users who can override lock down mode and access the DCUI (DCUI.Access)
config|security|shell_interactive_timeout	Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut)	Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut)
config|security|shell_timeout	Shell timeout (UserVars.ESXiShellTimeOut)	Shell timeout (UserVars.ESXiShellTimeOut)

Table 7-128. Configuration Properties Collected for Host System Objects (continued)

Property Key	Property Name	Description
config|security|dvfilter_bind_address	Dvfilter bind ip address (Net.DVFilterBindIpAddress)	Dvfilter bind ip address (Net.DVFilterBindIpAddress)
config|security|syslog_dir	Log directory (Syslog.global.logDir)	Log directory (Syslog.global.logDir)
config|security|firewallRule|allowedHosts	Allowed hosts	Allowed hosts in the firewall configuration
config|security|service|isRunning	Running	Indicates whether a service is running or not. Services are: Direct Console UI, ESXi shell, SSH, or NTP Daemon.
config|security|service|ruleSet	Ruleset	Ruleset for each service.
config|security|service|policy	Policy	Policy for each service.

Note Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

For more information on the *vSphere Hardening Guide* alerts, see Customize a Policy to Enable the vSphere Hardening Guide Alerts.

Table 7-129. Hardware Properties Collected for Host System Objects

Property Key	Property Name	Description
hardware|memorySize	Memory Size	Memory Size
hardware|cpuInfo|numCpuCores	Number of CPU Cores	Number of CPU Cores
hardware|cpuInfo|hz	CPU Speed per Core	CPU Speed per Core
hardware|cpuInfo|numCpuPackages	Number of CPU Packages	Number of CPU Packages
hardware|cpuInfo|powerManagementPolicy	Active CPU Power Management Policy	Active CPU Power Management Policy
hardware|cpuInfo|powerManagementTechnology	Power Management Technology	Power Management Technology
hardware|cpuInfo|biosVersion	BIOS Version	BIOS Version

Table 7-130. Runtime Properties Collected for Host System Objects

Property Key	Property Name	Description
runtime|connectionState	Connection State	Connection State
runtime|powerState	Power State	Power State
runtime|maintenanceState	Maintenance State	Maintenance State
runtime|memoryCap	Memory Capacity	Memory Capacity

Table 7-131. Configuration Manager Properties Collected for Host System Objects

Property Key	Property Name	Description
configManager|memoryManager|consoleReservationInfo|serviceConsoleReserved	Service Console Reserved	Service console reserved memory

Table 7-132. CPU Usage Properties Collected for Host System Objects

Property Key	Property Name	Description
cpu|speed	CPU	CPU Speed
cpu|cpuModel	CPU Model	CPU Model

Table 7-133. Network Properties Collected for Host System Objects

Property Key	Property Name	Description
net|maxObservedKBps	Highest Observed Throughput	Highest Observed Throughput (KBps)
net|mgmt_address	Management Address	Management Address
net|ip_address	IP Address	IP Address
net|discoveryProtocol|cdp|managementIpAddress	Management IP Address	Management IP Address
net|discoveryProtocol|cdp|systemName	System Name	System Name
net|discoveryProtocol|cdp|portName	Port Name	Port Name
net|discoveryProtocol|cdp|vlan	VLAN	VLAN
net|discoveryProtocol|cdp|mtu	MTU	MTU
net|discoveryProtocol|cdp|hardwarePlatform	Hardware Platform	Hardware Platform
net|discoveryProtocol|cdp|softwareVersion	Software Version	Software Version
net|discoveryProtocol|cdp|timeToLive	Time to Live	Time to Live
net|discoveryProtocol|lldp|managementIpAddress	Management IP Address	Management IP Address
net|discoveryProtocol|lldp|systemName	System Name	System Name
net|discoveryProtocol|lldp|portName	Port Name	Port Name
net|discoveryProtocol|lldp|vlan	VLAN	VLAN
net|discoveryProtocol|lldp|timeToLive	Time to Live	Time to Live

Table 7-134. System Properties Collected for Host System Objects

Property Key	Property Name	Description
sys|build	Build number	VMWare build number
sys|productString	Product String	VMWare product string

Table 7-135. Summary Properties Collected for Host System Objects

Property Key	Property Name	Description
summary|version	Version	Version
summary|hostuuid	Host UUID	Host UUID
summary|evcMode	Current EVC Mode	Current EVC Mode
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name
summary|parentCluster	Parent Cluster	Parent Cluster
summary|parentDatacenter	Parent Datacenter	Parent Datacenter
summary|parentVcenter	Parent Vcenter	Parent Vcenter

Table 7-136. Datastore Properties Collected for Host System Objects

Property Key	Property Name	Description
datastore| maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore| maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Cluster Compute Resource Properties

vRealize Operations Manager collects configuration and summary properties for cluster compute resource objects.

Table 7-137. Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 7-138. Summary Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
summary|parentDatacenter	Parent Datacenter	Parent Datacenter
summary|parentVcenter	Parent Vcenter	Parent Vcenter

Table 7-138. Summary Properties Collected for Cluster Compute Resource Objects (continued)

Property Key	Property Name	Description
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Table 7-139. DR, DAS, and DPM Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
configuration|drsconfig|enabled	Enabled	Indicates whether DRS is enabled
configuration|drsconfig| defaultVmBehavior	Default DRS Behaviour	Default DRS Behaviour
configuration|drsconfig|affinityRules	Affinity Rules	DRS Affinity Rules
configuration|dasconfig|enabled	HA Enabled	HA Enabled
configuration|dasconfig| admissionControlEnabled	Admission Control Enabled	Admission Control Enabled
configuration|dpmconfiginfo|enabled	DPM Enabled	DPM Enabled
configuration|dpmconfiginfo| defaultDpmBehavior	Default DPM Behaviour	Default DPM Behaviour

DRS properties are collected for disaster recovery. DAS properties are collected for high availability service, formerly distributed availability service. DPM properties are collected for distributed power management.

Resource Pool Properties

vRealize Operations Manager collects configuration, CPU, memory, and summary properties for resource pool objects.

Table 7-140. Configuration Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
config|name	Name	Name
config|cpuAllocation|reservation	Reservation	CPU reservation
config|cpuAllocation|limit	Limit	CPU limit
config|cpuAllocation| expandableReservation	Expandable Reservation	CPU expandable reservation
config|cpuAllocation|shares|shares	Shares	CPU shares
config|memoryAllocation|reservation	Reservation	Memory reservation
config|memoryAllocation|limit	Limit	Memory limit
config|memoryAllocation| expandableReservation	Expandable Reservation	Memory expandable reservation
config|memoryAllocation|shares|shares	Shares	Memory shares

Table 7-141. CPU Usage Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
cpu|limit	CPU Limit	CPU Limit
cpu|reservation	CPU reservation	CPU Reservation
cpu|expandable_reservation	CPU expandable reservation	CPU Expandable Reservation
cpu|shares	CPU Shares	CPU Shares
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s)

Table 7-142. Memory Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
mem|limit	Memory limit	Memory limit
mem|reservation	Memory reservation	Memory reservation
mem|expandable_reservation	Memory expandable reservation	Memory expandable reservation
mem|shares	Memory Shares	Memory Shares

Table 7-143. Summary Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Data Center Properties

vRealize Operations Manager collects configuration and summary properties for data center objects.

Table 7-144. Configuration Properties Collected for Data Center Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 7-145. Summary Properties Collected for Data Center Objects

Property Key	Property Name	Description
summary|parentVcenter	Parent Vcenter	Parent Vcenter
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Storage Pod Properties

vRealize Operations Manager collects configuration and summary properties for storage pod objects.

Table 7-146. Configuration Properties Collected for Storage Pod Objects

Property Key	Property Name	Description
config|name	Name	Name
config|sdrsconfig| vmStorageAntiAffinityRules	VM storage antiaffinity rules	Storage Distributed Resource Scheduler (SDRS) VM anti-affinity rules
config|sdrsconfig|vmdkAntiAffinityRules	VMDK antiaffinity rules	Storage Distributed Resource Scheduler (SDRS) Virtual Machine Disk (VMDK) anti-affinity rules

VMware Distributed Virtual Switch Properties

vRealize Operations Manager collects configuration and summary properties for VMware distributed virtual switch objects.

Table 7-147. Configuration Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 7-148. Capability Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
capability|nicTeamingPolicy	NIC Teaming Policy	NIC Teaming Policy

Distributed Virtual Port Group Properties

vRealize Operations Manager collects configuration and summary properties for distributed virtual port group objects.

Table 7-149. Configuration Properties Collected for Distributed Virtual Port Group Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 7-150. Summary Properties Collected for Distributed Virtual Port Group Objects

Property Key	Property Name	Description
summary|active_uplink_ports	Active DV uplinks	Active DV uplinks

Datastore Properties

vRealize Operations Manager collects configuration, summary, and properties about datastore use for datastore objects.

Table 7-151. Configuration Properties Collected for Datastore Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 7-152. Summary Properties Collected for Datastore Objects

Property Key	Property Name	Description
summary|diskCapacity	Disk Capacity	Disk Capacity
summary|isLocal	Is Local	Is local datastore
summary|customTag|customTagValue	Value	Custom Tag Value
summary|accessible	Datastore Accessible	Datastore Accessible

Table 7-153. Datastore Properties Collected for Datastore Objects

Property Key	Property Name	Description
datastore|hostcount	Host Count	Host Count
datastore|hostScsiDiskPartition	Host SCSI Disk Partition	Host SCSI Disk Partition
datastore| maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore| maxObservedReadLatency	Highest Observed Read Latency	Highest Observed Read Latency
datastore| maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore| maxObservedWriteLatency	Highest Observed Write Latency	Highest Observed Write Latency
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Self-Monitoring Properties for vRealize Operations Manager

vRealize Operations Manager uses the vRealize Operations Manager adapter to collect properties that monitor its own objects. These self-monitoring properties are useful for monitoring changes within vRealize Operations Manager.

Analytics Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager analytics service.

Table 7-154. Properties Collected for Analytics Service Objects

Property Key	Property Name	Description
HAEnabled	HA Enabled	Indicates HA is enabled with a value of 1, disabled with a value of 0.
ControllerDBRole	Role	Indicates persistence service role for the controller: 0 – Master, 1 – Replica, 4 – Client..
ShardRedundancyLevel	Shard redundancy level	The target number of redundant copies for Object data.

Table 7-154. Properties Collected for Analytics Service Objects (continued)

Property Key	Property Name	Description
LocatorCount	Locator Count	The number of configured locators in the system
ServersCount	Servers Count	The number of configured servers in the system

Node Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager node objects.

Table 7-155. Configuration Properties Collected for Node Objects

Property Key	Property Name	Description
config|numCpu	Number of CPU	Number of CPUs
config|numCoresPerCpu	Number of cores per CPU	Number of cores per CPU
config|coreFrequency	Core Frequency	Core Frequency

Table 7-156. Memory Properties Collected for Node Objects

Property Key	Property Name	Description
mem|RAM	System RAM	System RAM

Table 7-157. Service Properties Collected for Node Objects

Property Key	Property Name	Description
service|proc|pid	Process ID	Process ID

Remote Collector Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager remote collector objects.

Table 7-158. Configuration Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
config|numCpu	Number of CPU	Number of CPUs
config|numCoresPerCpu	Number of cores per CPU	Number of cores per CPU
config|coreFrequency	Core Frequency	Core Frequency

Table 7-159. Memory Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
mem|RAM	System RAM	System RAM

Table 7-160. Service Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
service|proc|pid	Process ID	Process ID

Properties for vSAN

vRealize Operations Manager displays object properties for vSAN.

Properties for vSAN Disk Groups

The vRealize Operations Manager displays the following property for vSAN disk groups:

- vSAN Disk Groups:Configuration|vSAN Configuration

Properties for vSAN Cluster

The vRealize Operations Manager displays the following properties for vSAN cluster.

- Cluster Configuration|vSAN|Deduplication and Compression Enabled
- Cluster Configuration|vSAN|Preferred fault domain
- Cluster Configuration|vSAN|Stretched Cluster
- Cluster Configuration|vSAN|vSAN Configuration

Properties for vSAN Enabled Host

The vRealize Operations Manager displays the following property for vSAN enabled host.

- Configuration|vSAN Enabled

Properties for vSAN Cache Disk

The vRealize Operations Manager displays the following properties for vSAN cache disk.

Properties for vSAN include:

Component	Metrics
Configuration	■ Configuration Properties|Name ■ Configuration Properties|Size ■ Configuration Properties|Vendor ■ Configuration Properties|Type ■ Configuration Properties|Queue Depth
SCSI SMART Statistics	■ SCSI SMART Statistics|Media Wearout Indicator Threshold ■ SCSI SMART Statistics|Write Error Count Threshold ■ SCSI SMART Statistics|Read Error Count Threshold ■ SCSI SMART Statistics|Reallocated Sector Count Threshold ■ SCSI SMART Statistics|Raw Read Error Rate Threshold ■ SCSI SMART Statistics|Drive Temperature Threshold ■ SCSI SMART Statistics|Drive Rated Max Temperature Threshold ■ SCSI SMART Statistics|Write Sectors TOT Count Threshold ■ SCSI SMART Statistics|Read Sectors TOT Count Threshold ■ SCSI SMART Statistics|Initial Bad Block Count Threshold

Properties for vSAN Capacity Disk

The vRealize Operations Manager displays the following properties for vSAN capacity disk.

Properties for vSAN include:

Component	Metrics
Configuration	■ Configuration Properties|Name ■ Configuration Properties|Size ■ Configuration Properties|Vendor ■ Configuration Properties|Type ■ Configuration Properties|Queue Depth
SCSI SMART Statistics	■ SCSI SMART Statistics|Media Wearout Indicator Threshold ■ SCSI SMART Statistics|Write Error Count Threshold ■ SCSI SMART Statistics|Read Error Count Threshold ■ SCSI SMART Statistics|Reallocated Sector Count Threshold ■ SCSI SMART Statistics|Raw Read Error Rate Threshold ■ SCSI SMART Statistics|Drive Temperature Threshold ■ SCSI SMART Statistics|Drive Rated Max Temperature Threshold ■ SCSI SMART Statistics|Write Sectors TOT Count Threshold ■ SCSI SMART Statistics|Read Sectors TOT Count Threshold ■ SCSI SMART Statistics|Initial Bad Block Count Threshold