

Help

10 MAY 2022

vRealize Operations 8.3

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 © 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

VMware vRealize Operations Manager 8.3 Help 12

1 About VMware vRealize Operations Manager 13

2 Planning 14

Continuous Availability FAQs 14

Reference Architecture 19

Best Practices for Deploying vRealize Operations Manager 19

Initial Considerations for Deploying vRealize Operations Manager 22

Scalability Considerations 24

High Availability Considerations 27

Continuous Availability Considerations 27

Adapter and Management Packs Considerations 29

Hardware Requirements for Analytics Nodes, Witness Nodes, and Remote Collectors 30

Port Requirements for vRealize Operations Manager 31

Small Deployment Profile for vRealize Operations Manager 32

Medium Deployment Profile for vRealize Operations Manager 33

Large Deployment Profile for vRealize Operations Manager 35

Extra Large Deployment Profile for vRealize Operations Manager 37

Secure Configuration 40

vRealize Operations Manager Security Posture 40

Secure Deployment of vRealize Operations Manager 41

Secure Configuration of vRealize Operations Manager 42

Network Security and Secure Communication 72

Auditing and Logging on your vRealize Operations Manager System 104

3 Installing 106

About Installing 106

Workflow of vRealize Operations Manager Installation 106

Sizing the Cluster 108

Complexity of Your Environment 109

Cluster Nodes 111

About Remote Collector Nodes 112

About High Availability 113

About vRealize Operations Manager Continuous Availability 114

Preparing for Installation 116

Requirements 117

Installing vRealize Operations Manager 121

Deployment of vRealize Operations Manager	121
Installation Types	124
Installing vRealize Operations Manager on VMware Cloud on AWS	130
Installing vRealize Operations Manager for Azure VMware Solution	136
Installing vRealize Operations Manager for Google Cloud VMware Engine	140
Resize your Cluster by Adding Nodes	144
Gathering More Data by Adding a Remote Collector Node	145
Adding High Availability	147
Adding Continuous Availability	148
Cluster and Node Maintenance	150
Troubleshooting	155
Post-Installation Considerations	156
About Logging In	156
After You Log In	157
Secure the Console	158
Log in to a Remote Console Session	159
About New Installations	160
Upgrade, Backup and Restore	161
Obtain the Software Update PAK File	161
Create a Snapshot as Part of an Update	162
How To Preserve Customized Content	162
Back Up and Restore	163
Software Updates	164
Before Upgrading to vRealize Operations Manager 8.3	167

4 Configuring 171

Connecting to Data Sources	172
Solutions Repository	174
Managing Solutions in vRealize Operations Manager	175
Managing Solution Credentials	182
Managing Collector Groups	184
VMware vSphere Solution	187
VMware Cloud on AWS	197
Azure VMware Solution	201
Google Cloud VMware Engine	202
AWS	204
Microsoft Azure	216
Application Monitoring	222
Service Discovery	323
Log Insight	332
Business Management	336

vRealize Automation 8.X	361
vSAN	370
vRealize Network Insight	375
End Point Operations Management Solution	376
NSX-T	437
Configuring Alerts and Actions	438
Triggered Alerts	438
Types of Alerts	443
Alert Information	443
Configuring Alerts	445
Viewing Actions	515
Configuring Policies	525
Policies	525
Operational Policies	529
Types of Policies	530
Using the Policy Workspace to Create and Modify Operational Policies	533
Define Monitoring Goals for vRealize Operations Manager Solutions	552
Configuring Compliance	553
What Are Compliance Benchmarks	553
How To Configure Compliance Benchmarks	557
Configuring Super Metrics	560
Create a Super Metric	562
Enhancing Your Super Metrics	565
Exporting and Importing a Super Metric	566
Super Metrics Tab	567
Configuring Objects	573
Object Discovery	574
Configuring Data Display	607
Widgets	608
Dashboards	748
Views	760
Reports	780
Configuring Administration Settings	791
License Keys	791
License Groups	793
Maintenance Schedules	795
Manage Maintenance Schedules	796
Managing Users and Access Control	796
Passwords and Certificates	834
Modifying Global Settings	845
Managing Content	850

Transfer Ownership of Dashboards and Report Schedules	852
Logs	853
Create a Support Bundle	855
Dynamic Thresholds	857
Adapter Redescribe	857
Customizing Icons	858
Allocate More Virtual Memory	861
About the Administration Interface	861
Cluster Status and Management	861
Logs	864
Support Bundles	865
Update the Reference Database for vRealize Operations Manager	866
Enable FIPS - Admin UI	866
Configuring and Using Workload Optimization	867
Configuring Workload Optimization	868
Using Workload Optimization	872
Workload Optimization Page	876
Rightsizing	881
Manage Optimization Schedules	883
Workload Automation Policy Settings	884
View DRS Summary	884
Optimization Schedules	885
Optimize Placement	886

5 Predefined Dashboards 888

The Getting Started Page	891
Availability Dashboards	896
VM Availability Dashboard	896
vSphere Availability Dashboard	898
Ping Overview Dashboard	899
Capacity Dashboards	900
Cluster Capacity Dashboard	901
Datastore Capacity Dashboard	903
ESXi Capacity Dashboard	905
VM Capacity Dashboard	906
VM Reclamation Dashboard	907
vSAN Capacity Dashboard	908
vSAN Stretched Clusters	909
Configuration Dashboards	910
Cluster Configuration Dashboard	912
ESXi Configuration Dashboard	914

Network Configuration Dashboard	915
VM Configuration Dashboard	916
vSAN Configuration Dashboard	918
Workload Management Configuration Dashboard	919
Consumer \ Correct it? Dashboard	919
Consumer \ Optimize it? Dashboard	921
Consumer \ Simplify it?	922
Consumer \ Update it? Dashboard	924
Provider \ Correct it? Dashboard	924
Provider \ Optimize it? Dashboard	926
Provider \ Simplify it? Dashboard	927
Provider \ Update it? Dashboard	928
Cost Dashboards	929
Assess Cost Dashboards	930
Base Rate Analysis Dashboard	930
Datacenter Cost Drivers Dashboard	931
Showback Dashboard	931
Performance Dashboards	932
Guest OS Performance Profiling Dashboard	938
Network Top Talkers Dashboard	940
Storage Heavy Hitters Dashboard	941
VM Contention Dashboard	942
VM Utilization Dashboard	944
Troubleshoot an Application Dashboard	945
Cluster Contention Dashboard	946
Cluster Utilization Dashboard	949
VM Rightsizing Dashboard	950
Datastore Performance Dashboard	951
ESXi Contention Dashboard	953
ESXi Utilization Dashboard	954
Network Performance Dashboard	955
vSAN Contention Dashboard	956
vSAN Utilization Dashboard	959
vSAN File Services	959
Dashboard Library	960
Deprecated Dashboards	960
Executive Summary Dashboards	974
Network Operation Center	976
Software Defined Wide Area Network Dashboard	981
Troubleshoot SD-WAN Dashboard	981
Troubleshoot SD-WAN Gateway Dashboard	982

Troubleshoot SD-WAN Orchestrator Dashboard	982
vRealize Automation 8.x Dashboards	983
Cloud Automation Environment Overview	983
Cloud Automation Project Cost Overview	983
Cloud Automation Resource Consumption Overview	984
Cloud Automation Top-N Dashboard	984
Service Discovery Dashboards	985
Service Distribution Dashboard	985
Service Relationships Dashboard	985
Service Visibility Dashboard	986
Virtual Machine Relationships Dashboard	986
Inventory Dashboards	986
vSphere Compute Inventory Dashboard	987
vSphere Network Inventory Dashboard	987
vSphere Storage Inventory Dashboard	988
Workload Management Inventory Dashboard	988
Microsoft Azure Dashboards	989
AWS Dashboards	990
AWS Instance Utilization Dashboard	992
AWS Auto Scaling Group Dashboard	992
AWS Troubleshooting Dashboard	992
AWS Instance Heatmap Dashboard	993
AWS Volume Performance Dashboard	993
AWS Disk Space Dashboard	993
Alerts	993
Dashboards in VMware Cloud on AWS	993
VMC Capacity Dashboard	993
VMC Cost Overview Dashboard	994
VMC Inventory Dashboard	994
VMC Management VM Monitoring Dashboard	995
VMC Utilization and Performance Dashboard	995
VMC Configuration Maximums Dashboard	996
Dashboards in NSX-T Management Pack	997
NSX-T Configmax Metrics	997

6 Monitoring Objects in Your Managed Environment 999

Enhanced Search Capability	999
What to Do When...	1001
User Scenario: A User Calls with a Problem	1001
User Scenario: An Alert Arrives in Your Inbox	1006
User Scenario: You See Problems as You Monitor the State of Your Objects	1015

Troubleshooting Workbench Home Page	1025
Discovering Potential Evidences Using the Troubleshooting Workbench	1025
Monitoring and Responding to Alerts	1027
Monitoring Alerts	1027
Monitoring and Responding to Problems	1032
Evaluating Object Information Using Badge Alerts and the Summary Tab	1033
Investigating Object Alerts	1057
Evaluating Metric Information	1065
Capacity Tab Overview	1072
Using Troubleshooting Tools to Resolve Problems	1074
Creating and Using Object Details	1081
Examining Relationships in Your Environment	1093
User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options	1095
Running Actions from vRealize Operations Manager	1100
Run Actions from Toolbars in vRealize Operations Manager	1100
Troubleshoot Actions in vRealize Operations Manager	1131
Monitor Recent Task Status	1133
Troubleshoot Failed Tasks	1138
Viewing Your Inventory	1146
Inventory Tab	1146

7 Capacity Optimization for Your Managed Environment 1147

Capacity Analytics	1148
Example: Excluding VMs from Reclaim Action	1153
What-If Analysis: Modeling Workload, Capacity, or Migration Planning	1154
Example: Run a What-If Scenario	1155
Example: Import Workload from an Existing VM Scenario	1156
Allocation Model	1158
Capacity Overview	1158
Reclaim	1162
Reclamation Settings	1166
What-If Analysis - Workload Planning: Traditional	1167
Add or Remove VMs	1170
What-If Analysis - Infrastructure Planning: Traditional	1173
Add or Remove Hosts	1174
What-If Analysis - Workload Planning: Hyperconverged	1175
Add or Remove VMS	1175
What-If-Analysis - Infrastructure Planning: Hyperconverged	1177
Add or Remove HCI Nodes	1178
What-If-Analysis - Migration Planning: Public Cloud	1179
Migration Planning	1180

What-If Analysis - Data Center Comparison	1182
Datacenter Comparison	1183
Custom Profiles	1184
Custom Profiles Details and Related Policies	1184
Custom Profiles Add and Edit Workspace	1185
Custom Data Centers in vRealize Operations Manager	1185
Custom Datacenters List	1186
Custom Datacenters Add and Edit Workspace	1187

8 Metric, Property, and Alert Definitions 1188

Metric Definitions in vRealize Operations Manager	1188
Metrics for vCenter Server Components	1188
Operating System Metrics	1310
Application Service Metrics	1314
VeloCloud Application Service Metrics	1353
Remote Check Metrics	1357
Service Discovery Metrics	1358
Calculated Metrics	1359
Self-Monitoring Metrics for vRealize Operations Manager	1371
vRealize Automation 8.x Metrics	1402
Metrics for vSAN	1405
Metrics for the Operating Systems and Remote Service Monitoring Plug-ins in End Point Operations Management	1417
Metrics for Microsoft Azure	1439
Metrics for Management Pack for AWS	1448
Metrics in VMware Cloud on AWS	1473
Metrics in NSX-T Adapter	1480
Alert Definitions in vRealize Operations Manager	1488
Cluster Compute Resource Alert Definitions	1489
Host System Alert Definitions	1494
vRealize Automation Alert Definitions	1507
vSAN Alert Definitions	1508
Alerts in the vSphere Web Client	1521
vSphere Distributed Port Group	1521
Virtual Machine Alert Definitions	1522
vSphere Distributed Switch Alert Definitions	1529
vCenter Server Alert Definitions	1530
Datastore Alert Definitions	1532
Data Center Alert Definitions	1537
Custom Data Center Alert Definitions	1538
vSphere Pod Alert Definitions	1539
VMware Cloud on AWS Alert Definitions	1543

Property Definitions in vRealize Operations Manager	1546
Properties for vCenter Server Components	1547
Self-Monitoring Properties for vRealize Operations Manager	1578
Service Discovery Properties	1579
Properties for vSAN	1581
Properties for vRealize Automation 8.x	1583
Properties in the NSX-T Adapter	1584
Placement Group Properties	1589
Properties for VeloCloud Gateway	1590
Properties for VeloCloud Orchestrator	1590

VMware vRealize Operations Manager 8.3 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.

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

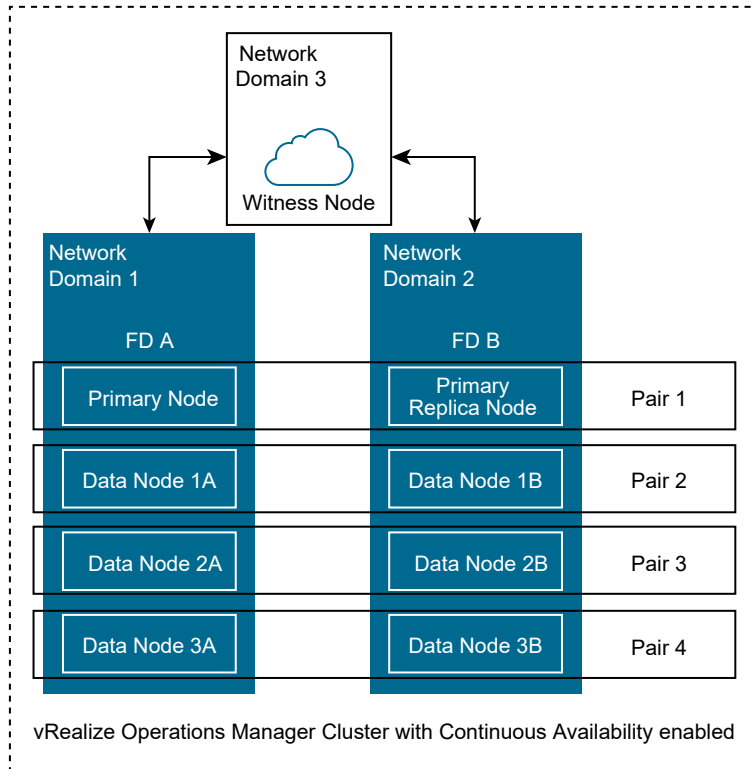
You plan your environment with recommendations for deployment and secure baseline for the deployment of vRealize Operations Manager.

This chapter includes the following topics:

- [Continuous Availability FAQs](#)
- [Reference Architecture](#)
- [Secure Configuration](#)

Continuous Availability FAQs

With the introduction of continuous availability in vRealize Operations Manager 8, there have been several frequently asked questions. This section is to help increase awareness and knowledge about continuous availability.



How is the data stored in analytics nodes?

When an object is discovered, vRealize Operations Manager determines which node to keep the data, then copies (duplicates) the data to its pair node in the other fault domain. Every object is stored in two analytics nodes (node pairs) across the fault domains and they are always synchronized.

As an example, vRealize Operations Manager has eight analytics nodes, CA is enabled, and as a result each fault domain has four analytics nodes (see above diagram).

When a new object is discovered, vRealize Operations Manager decides to store the data in “Data Node 2B” (primary) and automatically a copy of the data will be saved in “Data Node 2A” (secondary).

If somehow “FD A” becomes unavailable, then “primary” data from “Data Node 2B” will be used.

If somehow “FD B” becomes unavailable, then “secondary” data from “Data Node 2A” will be used.

Which situations break a continuous availability cluster? Simultaneously losing the primary node or primary replica node and data nodes, or two or more data nodes in both fault domains, are not supported.

Each analytics node from fault domain 1 has its node pair in fault domain 2 or vice versa.

Using the previously mentioned example, we will have four node pairs:

Primary + Replica Node

Data Node 1A (FD A) + Data Node 1B (FD B)

Data Node 2A (FD A) + Data Node 2B (FD B)

Data Node 3A (FD A) + Data Node 3B (FD B)

The two nodes of each node pair are always synchronized and storing the same data. Hence, the cluster continues to function without data loss while one node from all node pairs is available.

What happens if one data node from one of the fault domains becomes unavailable?

The cluster will be in a degraded state but continue to operate when one node becomes unavailable in either fault domain. There will be no data loss. The data node must be repaired or replaced so the cluster does not remain in a degraded state.

Will the cluster break if two data nodes in fault domain 1 and the primary replica node in fault domain 2 are lost?

In this example, the cluster will continue to work without data loss. If one analytics node from each node pair is still available, there will be no data loss.

What happens if an entire fault domain becomes unavailable?

The cluster will be in a degraded state but continue to operate when an entire fault domain becomes unavailable. There will be no data loss. The fault domain must be repaired and brought online so the cluster does not remain in a degraded state.

If the fault domain is unrecoverable, it is possible to replace the entire fault domain with newly deployed nodes. From the admin UI, only the primary replica node can be replaced. If the entire fault domain for the primary node is lost, you will need to wait until the primary node failover occurs and the primary replica node has been promoted as the new primary node.

What is the proper process to re-add a failed node to a fault domain? How long will it take to sync up?

The recommended procedure to re-add a failed node is to use the "Replace nodes of cluster" functionality within the admin UI. Once the replacement node has been added, the data will be synced. The sync time strongly depends on the object count, historical period of the objects, network bandwidth, and the load on the cluster.

What happens when network latency between fault domains exceeds 20 ms? How long can vRealize Operations Manager tolerate extended latency?

Adhering to latency requirements is necessary to achieve optimal performance. The latency between fault domains should be < 10 ms, with peaks up to 20 ms during 20 sec intervals. For more information about network latency guidelines, see the KB article [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783).

When network latency between fault domains goes above “20 ms during 20 sec intervals” for some period, but then recovers back to under 10 ms, how long does it take to resynchronize?

High latency does not mean that synchronization has stopped. When an object is discovered, vRealize Operations Manager will decide which node needs to keep the data (primary), then a second copy of the data will go to its node pair (secondary). Every object is stored in two analytics nodes (pairs) across both fault domains. Synchronization is an ongoing process where the secondary node is periodically syncs with the primary node. Synchronization is performed based on last synced timestamps of the primary and secondary nodes. Hence, there is no synchronization data queue in vRealize Operations Manager.

What is the actual witness node tolerance to missed polls?

Witness node operations are not poll based. The witness node interacts only when one of the nodes is not able to communicate (after various checks) with nodes from the other fault domain.

At what point in time will the primary node and primary replica node failover?

The failover occurs only when the primary node is no longer accessible or not alive.

When is the primary replica node promoted to the primary node?

The primary replica node is promoted to the primary node in only two cases:

- When the existing primary node is down.
- The associated fault domain is down/offline.

When the original primary node returns online, does it resume primary control? How does the data get synchronized?

When operations return to normal, with both primary node and primary replica node online, the newly promoted primary node (formerly primary replica node) remains the new primary node and the new primary replica (formerly primary node) gets synced with the new primary node.

What happens if connectivity between fault domains is completely interrupted, but then recovers?

If communications between the fault domains is completely interrupted for several minutes, then one of the fault domains will automatically go offline. After the network interruption is restored, admin user needs to manually bring the fault domain online which will begin the data synchronization.

What happens to the fault domains when the witness node becomes unavailable?

While both fault domains are healthy and communicating with each other, the unavailability of the witness node will have no effect on the cluster; vRealize Operations Manager will continue to function. If there is a communication issue between the fault domains, three situations could occur:

- Witness node is accessible from both fault domains – witness will bring one fault domain offline based on site health.
- Witness node is accessible from one fault domain only – The other fault domain will go offline automatically.
- Witness node is not accessible from both fault domains – Both fault domains will go offline.

When the offline fault domain becomes available again, will the fault domains synchronize all data collected during the communication outage?

The collected data is synchronized immediately once connectivity to the fault domain is restored and synchronized to capture all missed data.

What happens when an analytics node is not able to communicate to analytics nodes in the other fault domain?

If an analytics node is not able to communicate with all nodes from the other fault domain nor the witness node, it will go offline automatically. All nodes or entire fault domain that were taken offline automatically should be brought back online by the Admin user manually after ensuring that all communication issues have been resolved.

If the maximum number of nodes in a standard cluster is 8 extra-large nodes, which supports 320,000 objects, why is the maximum number of nodes in continuous availability more with 10 extra-large nodes, which supports 200,000 objects?

The 10 extra-large nodes are supported only in a continuous availability cluster and references a maximum of five extra-large nodes across two separate fault domains. This permits an increase to the number of nodes over a standard cluster and allows for collection for a greater number of objects.

A possible design is five extra-large nodes in fault domain 1, and 5 extra-large nodes in fault domain 2, with a witness node in a third site. The latency requirements must be met such that latency between fault domain 1 and fault domain 2 is <10 ms. Details about latency, packet loss and bandwidth are listed in the KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783).

Is a load balancer supported with Continuous Availability?

Yes, for more information about load balancer configuration, see vRealize Operations Manager Load Balancing Configuration guide available under Resources in the [vRealize Operations Manager Documentation page](#).

The documentation states, “When CA is enabled, the replica node can take over all functions that the primary node provides, in case of a primary node failure. The 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.”

During testing, by disconnecting the network interface on the primary node, the switchover to the new primary worked within 5 minutes, you get kicked out of the product UI or get strange errors.

The stated two or three minutes are approximate medium values, so 5 minutes is acceptable.

When the primary node is connected to the network again after a failover, what is the recommended procedure to return the original primary node to the primary role?

It is not necessary to roll back the primary replica node to the primary node role or vice versa. If you still want to restore the old primary node to the primary role, then use “Take Node Offline/Online” on the new primary node or its fault domain (where the original primary node resides)

Anytime a node goes offline or gets rebooted, is it necessary to bring the corresponding fault domain offline and then online to bring the node back online?

All nodes, after reboot or bringing it offline/online, will automatically continue to work. No additional steps are necessary.

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

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

Note The master node is now referred to as the primary node. The master replica node is now referred to as the primary replica node.

- Deploy analytics nodes in the same vSphere Cluster except when enabling Continuous Availability.
- Deploy analytics nodes with the same disk size on storage of the same type.
- When enabling Continuous Availability, separate analytics nodes into fault domains based on their physical location.
- 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 the resource reservation to ensure optimal performance. Ensure that the virtual CPU to physical CPU ratio is not negatively impacting the performance of analytics 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 the least number of nodes. For example, if the number of collecting is 120,000, configure the cluster with four extra-large nodes instead of 12 large nodes.
- Deploy an extra even number of nodes to enable Continuous Availability. If the current configuration is an odd number of analytics nodes, deploy an extra analytics node to create an even pairing.

Remote Collector Nodes

Remote collector nodes are additional cluster nodes that allow vRealize Operations Manager to gather more objects into its inventory for monitoring.

- Deploy remote collector nodes when the cluster is online.
- Deploy remote collector nodes one at a time. Adding multiple remote collectors in parallel can cause the cluster to crash.

Witness Nodes

A witness node is required when continuous availability is enabled to manage the analytics nodes in the fault domains.

- Deploy the witness node before enabling continuous availability.
- Deploy the witness node using the witness configuration.
- Deploy the witness node in a different cluster separate from the analytics 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.

- Utilize remote collector groups to separate data collection into fault domains when continuous availability is enabled.

vRealize Application Remote Collector and Telegraf Agents

- Deploy vRealize Application Remote Collector in the same vCenter Server as the end point VMs on which you want to deploy the Telegraf agents.

- Ensure that your operating system platform is supported by vRealize Application Remote Collector, and the most recent versions of Windows and Linux OS are supported.
- System times must be synchronized between vRealize Application Remote Collector, end point VMs, the vCenter Server, ESX host, and vRealize Operations Manager. To ensure synchronized time, use Network Time Protocol (NTP).
- Disable UAC on Endpoint VMs before installing the Telegraf agent. If you cannot do this due to security restrictions, see [KB article 70780](#) for a work around script.
- Ensure that the latest version of VMware Tools is installed on the end point VM on which you want to deploy the Telegraf agent.
- To deploy Telegraf agents onto end point VMs, ensure that the following prerequisites are met for the user account being used for deployment:

Windows - The user account must be either:

- An administrator account
- A non-administrator account that is a member of the built-in administrator group

Linux - The user account must be either:

- A root user with all privileges
- A non-root user with all privileges
- A non-root user with specific privileges

For more information, see User Account Prerequisites in the *vRealize Operations Manager Configuration Guide*.

Deployment Formats

Deploy vRealize Operations Manager with the same vRealize Operations Manager vApp version for the following node types:

- Primary
- Primary Replica
- Data
- Remote Collector
- Witness

See the *vRealize Operations Manager vApp Deployment and Configuration Guide* for more information.

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 320,000 monitored resources spread across eight extra-large analytics nodes.

Size your vRealize Operations Manager instance to ensure performance and support. For more information about sizing, refer to the KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783) .

Environment

Deploy analytics nodes in the same vSphere cluster and use identical or similar hosts and storage. If you cannot deploy analytics nodes in the same vSphere cluster, you must deploy them in the same geographical location.

When continuous availability is enabled, deploy analytics nodes in fault domains in the same vSphere cluster and use identical or similar hosts and storage. Fault domains are supported on vSphere stretched clusters.

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

- vMotion
- Storage vMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

Due to a high level of traffic between analytics nodes, all analytics nodes must be on the same VLAN and IP subnet, and that VLAN is not stretched between data centers, when continuous availability is not enabled.

When continuous availability is enabled, analytics nodes in fault domains should be located on the same VLAN and IP subnet, and communication between fault domains must be available. The witness node might be located in a separate VLAN and IP subnet but must be able to communicate with all analytics nodes.

Latency between analytics nodes cannot exceed 5 milliseconds, except when continuous availability is enabled, where latency between fault domains cannot exceed 10 milliseconds but analytics nodes, within each fault domain, still cannot exceed 5 milliseconds. The bandwidth must be equal to or faster than 10 GB per second.

If you deploy analytics nodes into a highly consolidated vSphere cluster, configure resource reservations. A full analytics node, for example a large analytics node that monitors 20,000 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 ratio is 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 and the witness node behind a firewall. You cannot use NAT between remote collectors or the witness node and analytics nodes.

Multiple Data Centers

vRealize Operations Manager can be stretched across data centers only when continuous availability is enabled. The fault domains may reside in separate vSphere clusters; however, all analytics nodes across both fault domains must reside in the same geographical location.

For example, the first datacenter is located in Palo Alto but is configured in two different buildings or in different locations of the city (downtown and mid-town) will have latency that is less than 5 milliseconds. The second datacenter is located in Santa Clara so the latency between the two datacenters is greater than 5 milliseconds but less than 10 milliseconds. Refer to the KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783) for network requirements.

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 centers. 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 you monitor collections to validate that they are completing in less than five minutes. Check the KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783) for latency, bandwidth and sizing requirements. If all requirements are met and collections are still not completing within the default 5 minutes 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 nodes, remote collector nodes, witness nodes, 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 is modified. For more information about root and intermediate certificates, see Specify the End Point Operations Management Agent Setup Properties in the *VMware vRealize Operations Manager Configuration Guide*.

Adapters

It is recommended that you configure adapters to remote collectors in the same data center as the analytics cluster for large and extra-large deployment profiles. Configuring adapters to remote collectors improves performance by reducing load on the analytics node. As an example, you might decide to configure an adapter to remote collectors if the total resources on a given analytics node begin to degrade the node's performance. You might configure the adapter to a large remote collector with the appropriate capacity.

Configure adapters to remote collectors when the number of resources the adapters are monitoring exceeds the capacity of the associated analytics node.

vRealize Application Remote Collector

For the production instance of vRealize Application Remote Collector and Telegraf agents to function optimally, your environment must adhere to certain configurations. You must review these configurations before you start deploying vRealize Application Remote Collector and Telegraf agents.

Option	Configurations
Sizing	The vRealize Application Remote Collector supports up to a maximum of 10,000 Telegraf agents using a large vRealize Application Remote Collector. Size your vRealize Application Remote Collector instance to ensure optimal performance and support. For more information about sizing, see the KB article, vRealize Operations Manager Sizing Guidelines (KB 2093783) .
Environment	Deploy the vRealize Application Remote Collector in the same vCenter Server as the end point VMs where you want to deploy Telegraf agents. Latency between the vRealize Application Remote Collector and a vRealize Operations Manager node cannot exceed 10 milliseconds.

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 *Deploying the vCenter Server Appliance* in the *VMware vSphere Documentation*. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

Load Balancer

For more information about load balancer configuration, see the *vRealize Operations Manager Load Balancing Guide*.

Scalability Considerations

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

Analytics Nodes

Analytics nodes consist of a primary node, a primary replica node, and data nodes.

For enterprise deployments of vRealize Operations Manager, deploy all nodes as medium, large or extra-large deployments, depending on sizing requirements and 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. It is recommended to scale up the analytics nodes in the cluster before scaling out the cluster with additional nodes. 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 the topic, *Add Data Disk Space to a vRealize Operations Manager vApp Node*. You cannot modify the disks of virtual machines that have a snapshot. You must remove all snapshots before you increase the disk size.

Scaling Horizontally (Adding nodes)

vRealize Operations Manager supports up to eight extra-large analytics nodes in a cluster, or up to 10 extra-large nodes in a cluster when continuous availability is enabled.

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

Witness Node

vRealize Operations Manager provides a single size regardless of the cluster size since the witness node does not collect nor process data.

Table 2-2. Witness Node Deployment Sizes

Node Size	vCPU	Memory
Witness	2	8 GB

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 2-3. Supported Remote Collector Sizes

Collector Size	Resources	End Point Operations Management Agents
Standard	6000	250
Large	32,000	2,000

For more information about sizing, see the KB article [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783).

vRealize Application Remote Collector

vRealize Operations Manager supports three sizes for application remote collectors; small, medium, and large. The number of Telegraf agents you want to deploy determines the size of vRealize Application Remote Collector you deploy.

Currently, vRealize Application Remote Collector can collect data on 20 different application sources.

If you have more than 6,000 Telegraf agents installed, increase vCPU and memory of the large configurations so that you can monitor up to 10,000 Telegraf agents.

An increase of memory usage depends on the number of services and their configurations on the VMs that are being monitored. When you monitor 1000 operating system objects, memory usage increases by around 1-1.5 GB.

Table 2-4. Supported vRealize Application Remote Collector Sizes

vRealize Application Remote Collector Size	Maximum number of Telegraf Agents Supported
Small	500
Medium	3000
Large	6000

High Availability Considerations

High availability creates a replica for the vRealize Operations Manager primary node and protects the analytics cluster against the loss of a node.

Cluster Management

Clusters consist of a primary node, a primary replica node, data nodes, and remote collector nodes.

Enabling High Availability within vRealize Operations Manager is not a disaster recovery solution. When you enable High Availability, information is stored (duplicated) in two different analytics nodes within the cluster. This doubles the system's compute and capacity requirements. If either the primary node or the primary replica node is permanently lost, then you must disable, and then re-enable High Availability to reassign the primary replica role to an existing node. This process, which includes a hidden cluster rebalance, can take a long time.

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you enable High Availability, you protect vRealize Operations Manager from data loss when only a single node is lost. If two or more nodes are lost, there may be permanent data loss. Deploy each analytics node to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the vRealize Operations Manager nodes remain on separate hosts.

Remote Collectors

In vRealize Operations Manager, you can create a collector group. A collector group is a collection of nodes (analytics nodes and remote collectors). You can assign adapters to a 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. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see [Adapter and Management Packs Considerations](#).

Continuous Availability Considerations

Continuous Availability (CA) separates the vRealize Operations Manager cluster into two fault domains and protects the analytics cluster against the loss of a fault domain.

Cluster Management

Clusters consist of a primary node, a primary replica node, a witness node, data nodes, and remote collector nodes.

Enabling Continuous Availability within vRealize Operations Manager is not a disaster recovery solution.

When you enable Continuous Availability, information is stored (duplicated) in two different analytics nodes within the cluster but stretched across fault domains. Due to sizing requirements, continuous availability requires doubling the system's compute and capacity requirements.

If either the primary node or primary replica node is permanently lost, then you must replace the lost node, which will become the new primary replica node. If it is necessary to have the new primary replica node as the primary node, then you can take the current primary node offline and wait until the primary replica node is promoted to the new primary node. Then bring the former primary node back online and it will be the new primary replica node.

Fault Domains

Fault domains consist of analytics nodes, separated into two zones.

A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. When configured, two fault domains enable vRealize Operations Manager to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

Witness Node

Witness node is a member of the cluster but not part of the analytics nodes.

To enable CA within vRealize Operations Manager, deploy the witness node in the cluster. The witness node does not collect nor store data.

The witness node serves as a tiebreaker when a decision must be made regarding availability of vRealize Operations Manager when the network connection between the two fault domains is lost.

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you enable continuous availability, you protect vRealize Operations Manager from data loss if an entire fault domain is lost. If node pairs are lost across fault domains, there may be permanent data loss.

Deploy analytics nodes, within each fault domain, to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the vRealize Operations Manager nodes remain on separate hosts.

Remote Collectors

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

When enabling continuous availability, collector groups can be created to collect data from adapters within each fault domain.

Collector groups do not have any correlation with fault domains. The functionality of a collector group is to collect data and provide it to the analytics nodes, which then vRealize Operations Manager decides how to keep the data.

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

Theoretically, you can install collectors in any place, provided the networking requirements are being met. However, from a failover perspective, it is not recommended to put all the collectors within a single fault domain. If all the collectors are directed to a single fault domain, vRealize Operations Manager stops receiving data if a network outage occurs affecting that fault domain.

The recommendation is to keep remote collectors outside of fault domains or keep half of the remote collectors in fault domain 1 and the remaining remote collectors in fault domain 2.

Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see [Adapter and Management Packs Considerations](#).

Adapter and Management Packs Considerations

Adapters and management packs have specific configuration considerations.

Normal Adapters

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

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

- VMware vSphere
- Management Pack for NSX for vSphere
- Management Pack for VMware Integrated OpenStack
- Management Pack for Storage Devices
- Management Pack for Log Insight

Hybrid Adapters

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

You must deploy hybrid adapters to a dedicated remote collector. Configure only one hybrid adapter type for each remote collector. 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 and extra-large analytics nodes can support 2,500 endpoint agents and large remote collectors can support 2,000 per node. To reduce the ingestion load on the cluster, you can point End Point Operations Management adapters at remote collectors. 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.

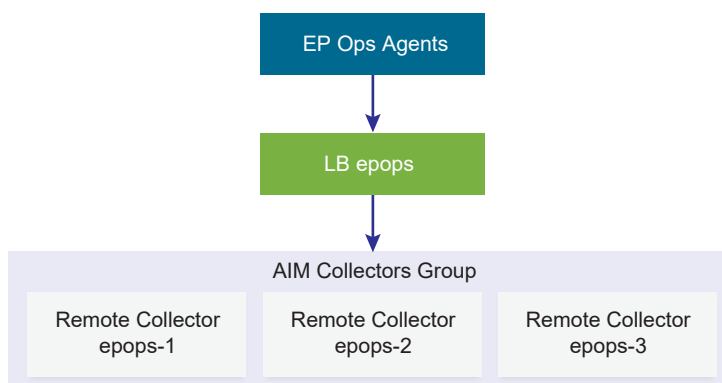
To reduce the cost of reconfiguring the system, it is recommended that you install End Point Operations Management agents against a DNS entry specific to End Point Operations Management agents if you plan to scale the system beyond a single node.

vRealize Application Management Pack

When you activate application monitoring in vRealize Operations Manager, as part of the process, you have to either download the vRealize Application Remote Collector OVA from within vRealize Operations Manager or externally from My VMware.

This is a dedicated virtual appliance which acts as a proxy between vRealize Operations Manager, the target vCenter Server, and the end point VMs where Telegraf agents are deployed.

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



Hardware Requirements for Analytics Nodes, Witness Nodes, and Remote Collectors

Analytics nodes, witness 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 2-5. Hardware Requirements for System Components

Server Roles	Virtual CPU	Memory	Storage Requirements
Small Analytics Node	4 vCPU	16 GB	1276 IOPS
Medium Analytics Node	8 vCPU	32 GB	1875 IOPS
Large Analytics Node	16 vCPU	48 GB	3750 IOPS
Extra Large Analytics Node	24 vCPU	128 GB	12758 IOPS
Standard Remote Collector	2 vCPU	4 GB	N/A
Large Remote Collector	4 vCPU	16 GB	N/A
Witness Node	2 vCPU	8 GB	N/A
Small vRealize Application Remote Collector	4 vCPU	8 GB	N/A
Medium vRealize Application Remote Collector	8 vCPU	16 GB	N/A
Large vRealize Application Remote Collector	16 vCPU	24 GB	N/A


CPU requirements are 2.0 GHz minimum. 2.4 GHz is recommended. 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. The cluster will perform better if the nodes stay within a single socket.

Port Requirements for vRealize Operations Manager

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

Port Requirements for vRealize Operations Manager



Ports Information for vRealize Operations Manager

Ports information for vRealize Operations Manager is available on [Ports and Protocol](#).

Small Deployment Profile for vRealize Operations Manager

The small deployment profile is intended for systems that manage up to 20,000 resources.

Virtual Appliance Name

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

Deployment Profile Support

The small deployment profile supports the following configuration.

- 20,000 resources
- 2,500 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 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 -> analytics-1.ra.local

Certificate

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

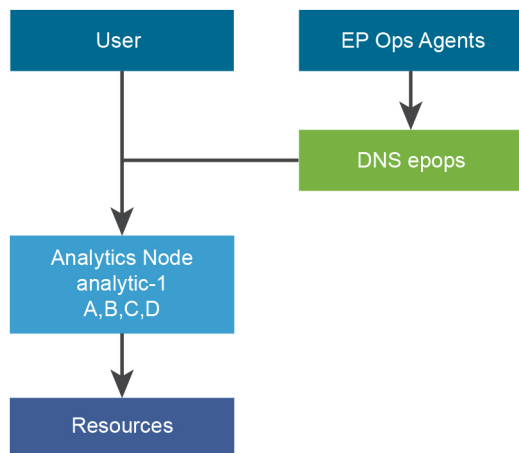
- DNS Name = *epops.refarch.local*
- DNS Name = *analytics-1.ra.local*

This is an example of a small deployment profile.

Table 2-6. Adapter Properties

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytics-1	A	2,000
DEFAULT	analytics-1	B	4,000
DEFAULT	analytics-1	C	2,000
DEFAULT	analytics-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 68,000 resources, 34,000 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.

- analytics-1.ra.lcoal

- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal
- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

Deployment Profile Support

The medium deployment profile supports the following configuration.

- 68,000 total resources, 34,000 enabled for HA
- 9,600 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 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 = *analytics-1.ra.local*

This is an example of a medium deployment profile.

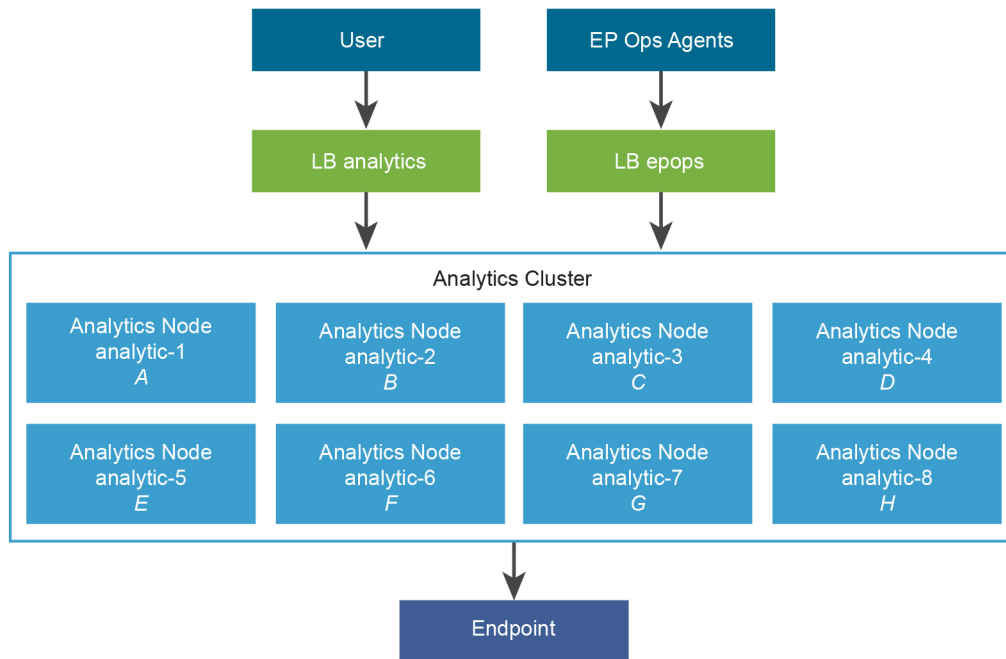
Table 2-7. Adapter Properties

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

Table 2-7. Adapter Properties (continued)

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytics-7	G	1,500
DEFAULT	analytics-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 128,000 resources, 64,000 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.

In addition, vRealize Application Remote Collector can be deployed to collect application level data for up to 6,000 end point VMs using Telegraf agents.

Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, large remote collectors for adapters, and large remote collectors for Telegraf agents.

- analytics-1.ra.lcoal
- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal

- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

Deployment Profile Support

The large deployment profile supports the following configuration.

- 128,000 total resources, 64,000 enabled for HA
- 6,000 Telegraf agents
- 20,000 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 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 = *analytics.refarch.local*
- DNS Name = *epops.refarch.local*
- DNS Name = *analytics-1.ra.local* to DNS Name = *analytics-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 2-8. 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

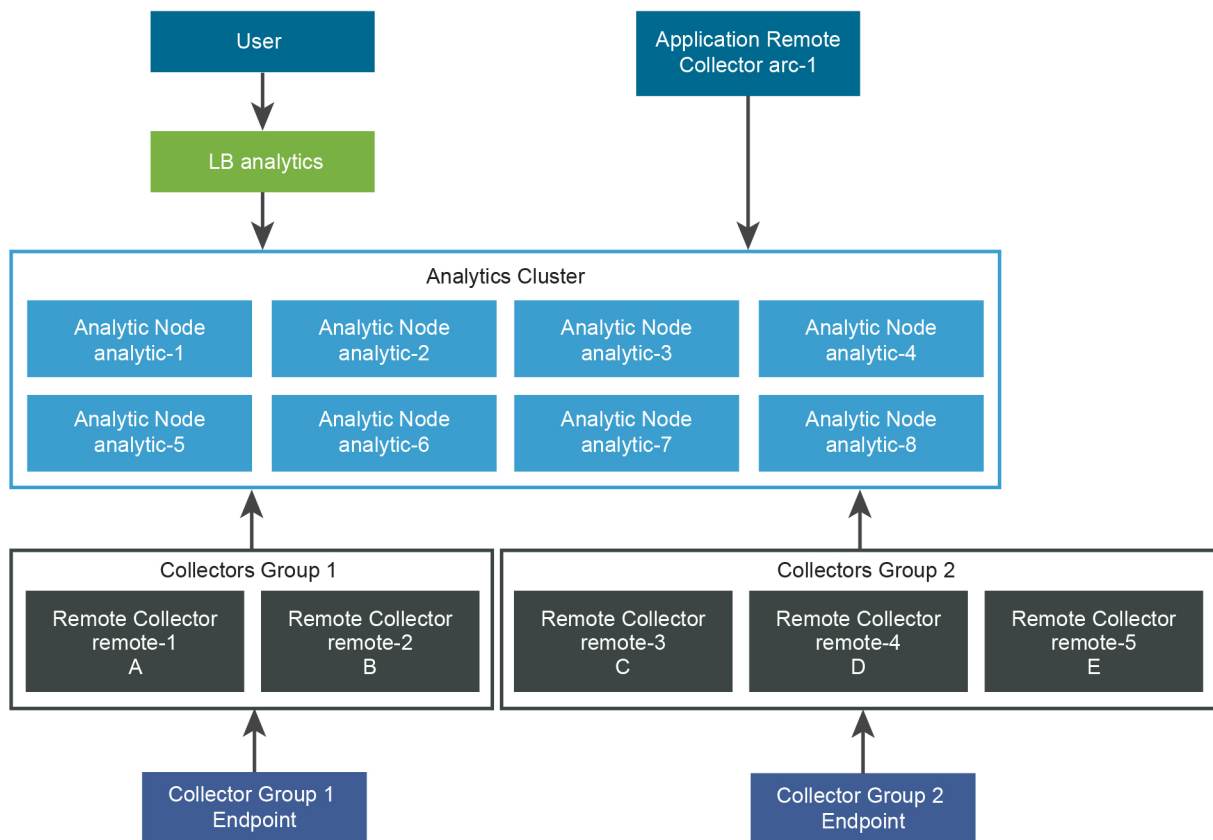
Table 2-8. Adapter Properties (continued)

Collector Group	Remote Collector	Adapter	Resources	End Point Operations Management Agents
		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 32,000 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 240,000 resources, 120,000 of which are enabled for Continuous 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 six extra-large analytics nodes. Large remote collectors for adapters, large remote collectors for End Point Operations Management agents, and witness node for continuous availability.

- `analytics-1.ra.local`
- `analytics-2.ra.local`
- `analytics-3.ra.local`
- `analytics-4.ra.local`
- `analytics-5.ra.local`
- `analytics-6.ra.local`
- `witness-1.ra.local`

Deployment Profile Support

- 240,000 total resources, 120,000 enabled for CA
- 20,000 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 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 = *analytics.refarch.local*
- DNS Name = *epops-a.refarch.local*
- DNS Name = *epops-b.refarch.local*
- DNS Name = *analytics-1.ra.local* to *analytics-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*
- DNS Name = *witness-1.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 2-9. 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
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

- 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 MD5/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/SHA256 hash output with the value posted on the VMware website. SHA256, SHA1, or MD5 hash should match.

Note The vRealize Operations Manager 6.x-x.pak/7.x-x.pak/8.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.

Enabling FIPS 140-2

FIPS 140-2 accreditation validates that an encryption solution meets a specific set of requirements designed to protect the cryptographic module from being cracked, altered, or otherwise tampered with. When FIPS 140-2 mode is enabled, any secure communication to or from vRealize Operations Manager 8.4 uses cryptographic algorithms or protocols that are allowed by the United States Federal Information Processing Standards (FIPS). FIPS mode turns on the cipher suites that comply with FIPS 140-2. Security related libraries that are shipped with vRealize Operations Manager 8.4 are FIPS 140-2 certified. However, the FIPS 140-2 mode is not enabled by default. FIPS 140-2 mode can be enabled if there is a security compliance requirement to use FIPS certified cryptographic algorithms with the FIPS mode enabled.

Note Enabling FIPS is a one-way action, and cannot be disabled after it is enabled.

Enable FIPS during the initial cluster deployment

- Ensure a new deployment of a vRealize Operations Manager cluster.
- Ensure that the Enable FIPS flag is appropriately used during the deployment of cluster nodes (OVF/OVA).

Enable FIPS on a working cluster

- 1 Navigate to `https://<VROPS IP>/admin/index.action`.
- 2 Login as an admin user.
- 3 Take the cluster offline to activate the Enable FIPS button in the **Administrator Settings** page.
- 4 Open the **Administrator Settings** tab in the left panel.
- 5 Click Enable FIPS under the **FIPS Setting** section.
- 6 Bring the cluster online.

Verify that FIPS mode is Enabled

From the Admin user interface:

- 1 Navigate to `https://<VROPS IP>/admin/index.action`.
- 2 Login as the admin user.
- 3 Open the **Administrator Settings** tab from the left panel.
- 4 A **FIPS 140-2 Status** message appears.

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 prompt for a new password, enter the root password that you want and note it for future reference.

- 4 Reenter the root password.

- 5 Log out of the console.

Change the Root Password

You can change the root password for any vRealize Operations Manager primary 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/system-password`. All hardened appliances enable `enforce_for_root` for the `pw_history` module, found in the `/etc/pam.d/system-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, the root account is set to a 365-day password expiry.

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 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 `#systemctl is-enabled sshd` command.
- 4 If the `sshd` service is disabled, run the `#systemctl enable sshd` command.
- 5 Run the `# systemctl start sshd` command to start the `sshd` service.
- 6 Run the `# systemctl stop sshd` command to stop the `sshd` service.

You can also enable or disable Secure Shell from the **SSH Status** column of the vRealize Operations Manager administration interface.

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 username -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 `ttty1` entry with `console`.

Restrict Secure Shell Access

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

Procedure

- 1 Open the `/etc/ssh/sshd_config` file on your virtual appliance host machine in a text editor.
- 2 Change the generic entry for your production environment to include only the local host entries and the management network subnet for secure operations.

Add the following line to the configuration file:

```
AllowUsers root@127.0.0.1 root@::1 root@10.0.0.*
```

In this example, all local host connections and connections that the clients make from the 10.0.0.0/24 subnet are allowed.

- 3 Save the file and close it.
- 4 Restart the SSH service by `systemctl restart sshd`.

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	aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-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 use 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
Network Sessions	MaxSessions 1
Strict Mode Checking	Strict Modes yes

Setting	Status
Privilege Separation	UsePrivilegeSeparation yes
rhosts RSA Authentication	RhostsRSAAuthentication no
Compression	Compression delayed or Compression no
Message Authentication code	hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha1-etm@openssh.com,hmac-sha2-512,hmac-sha2-256,hmac-sha1
User Access Restriction	PermitUserEnvironment no
KexAlgorithms	diffie-hellman-group14-sha1,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521

- 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
Ciphers	aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr
Message Authentication Codes	hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha1-etm@openssh.com,hmac-sha2-512,hmac-sha2-256,hmac-sha1

- 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 in the `/boot/grub/grub.cfg` file on your virtual appliances.
- 2 If no password exists, run the `/usr/bin/grub2-mkpasswd-pbkdf2` command on your virtual appliance.

A password is generated, and the command supplies the hash output.

- 3 Add following lines at the end of `/etc/grub.d/40_custom`.

```
set superusers="root"

password_pbkdf2 root <hash of password>
```

- 4 Backup `/boot/grub/grub.cfg` file by using:

```
cp /boot/grub/grub.cfg /boot/grub/grub.cfg.vropsbackup
```

- 5 Update the grub configuration by running the `/usr/sbin/grub2-mkconfig -o /boot/grub/grub.cfg` command.

What to do next

Note Important: Follow the upgrade procedure described below as otherwise, after upgrade, vRealize Operations Manager will not start.

Upgrade procedure for vRealize Operations Manager with a password protected boot loader.

- 1 Restore the old `grub.cfg` by running the following command:

```
cp /boot/grub/grub.cfg.vropsbackup /boot/grub/grub.cfg
```

- 2 Upgrade vRealize Operations Manager.
- 3 Follow all the steps described under **Set Boot Loader Authentication** after the upgrade of vRealize Operations Manager.

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:

```
root:x:0:0:root:/root:/bin/bash
bin:x:1:1:bin:/dev/null:/bin/false
daemon:x:6:6:Daemon User:/dev/null:/bin/false
messagebus:x:18:18:D-Bus Message Daemon User:/var/run/dbus:/bin/false
systemd-bus-proxy:x:72:72:systemd Bus Proxy:/:/bin/false
systemd-journal-gateway:x:73:73:systemd Journal Gateway:/:/bin/false
systemd-journal-remote:x:74:74:systemd Journal Remote:/:/bin/false
systemd-journal-upload:x:75:75:systemd Journal Upload:/:/bin/false
systemd-network:x:76:76:systemd Network Management:/:/bin/false
systemd-resolve:x:77:77:systemd Resolver:/:/bin/false
systemd-timesync:x:78:78:systemd Time Synchronization:/:/bin/false
nobody:x:65534:65533:Unprivileged User:/dev/null:/bin/false
sshd:x:50:50:sshd PrivSep:/var/lib/ssh:/bin/false
apache:x:25:25:Apache Server:/srv/www:/bin/false
```

```
ntp:x:87:87:Network Time Protocol:/var/lib/ntp:/bin/false
named:x:999:999::/var/lib/bind:/bin/false
admin:x:1000:1003::/home/admin:/bin/bash
postgres:x:1001:100::/var/vmware/vpostgres/9.6:/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.

```
root:x:0:admin
bin:x:1:daemon
sys:x:2:
kmem:x:3:
tape:x:4:
tty:x:5:
daemon:x:6:
floppy:x:7:
disk:x:8:
dialout:x:10:
audio:x:11:
video:x:12:
utmp:x:13:
usb:x:14:
cdrom:x:15:
adm:x:16:
messagebus:x:18:
systemd-journal:x:23:
input:x:24:
mail:x:34:
lock:x:54:
dip:x:30:
systemd-bus-proxy:x:72:
systemd-journal-gateway:x:73:
systemd-journal-remote:x:74:
systemd-journal-upload:x:75:
systemd-network:x:76:
systemd-resolve:x:77:
systemd-timesync:x:78:
nogroup:x:65533:
users:x:100:
sudo:x:27:
wheel:x:28:root,admin
sshd:x:50:
apache:x:25:admin,apache
ntp:x:87:
named:x:999:
vami:x:1000:root
admin:x:1003:
```


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 primary 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 -4 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.

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, TLS 1.0 and TLS 1.1 have also been disabled and only TLS 1.2 is enabled by default.

Note When you upgrade from vRealize Operations Manager 7.5 and above to 8.3, the user modifications to TLS settings are preserved. When you upgrade your vRealize Operations Manager instance from 7.0 to 8.3, both TLS 1.0 and TLS 1.1 are disabled on all the vRealize Operations Manager nodes. TLS 1.2 is the only protocol that is supported by default.

Verify the Correct Use of Protocols in Apache HTTPD

vRealize Operations Manager disables SSLv2, SSLv3, TLSv1, and TLSv1.1 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, SSLv3, TLSv1, and TLSv1.1 are disabled.

If the protocols are disabled, the command returns the following output: `SSLProtocol All -SSLv2 -SSLv3 -TLSv1 -TLSv1.1`.

- 2 To restart the Apache2 server, run the `systemctl restart httpd` command from the command prompt.

Verify the Correct Use of Protocols in the GemFire TLS Handler

vRealize Operations Manager disables SSLv3, TLS 1.0, and TLS 1.1 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:

```
1. # grep inter_cluster.supported_protocols /storage/vcops/user/conf/ssl/secure-communications.properties
or
2. # grep default.supported_protocols /storage/vcops/user/conf/ssl/secure-communications.properties
```

If the result of command 1 is blank, that means that the `inter_cluster` properties are not specified directly and it uses default values which you can obtain by command 2.

- 2 Re-enable TLS 1.0 and TLS 1.1.

- a Navigate to the administrator user interface to bring the cluster offline: `url/admin`.
- b Click **Bring Offline**.
- c To ensure that TLS 1.0 and TLS 1.1 are enabled, run the following commands:

If the result of command 1 is blank, use the following command:

```
sed -i "/^[^#]*default.supported_protocols/ c\default.supported_protocols = TLSv1.2 TLSv1.1 TLSv1" /storage/vcops/user/conf/ssl/secure-communications.properties
```

If the result of command 1 is not blank, use the following command:

```
sed -i "/^[^#]*inter_cluster.supported_protocols/ c\inter_cluster.supported_protocols = TLSv1.2 TLSv1.1 TLSv1" /storage/vcops/user/conf/ssl/secure-communications.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 HIGH:!aNULL!ADH:!EXP:!MD5:!3DES:!CAMELLIA:!PSK:!SRP:!
DH:@STRENGTH
```

- 2 To configure the correct use of cipher suites, run the `sed -i "/^[^#]*SSLCipherSuite/c\SSLCipherSuite HIGH:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP:\!DH:@STRENGTH"` /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 reenble DH, remove `!DH` from the cipher suites by running the `sed -i "/^[^#]*SSLCipherSuite/c\SSLCipherSuite HIGH:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP:@STRENGTH"` /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf command from the command prompt.
- 5 Run the `systemctl restart httpd` 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:

```
1. # grep inter_cluster.supported_cipher_suites /storage/vcops/user/conf/ssl/secure-
communications.properties
or
2. # grep default.supported_cipher_suites /storage/vcops/user/conf/ssl/secure-
communications.properties
```

If the result of command 1 is blank, that means that the `inter_cluster` properties are not specified directly and it uses default values which you can obtain by command 2.

The following result is expected:

```
inter_cluster. supported_cipher_suites =
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
```

If the result of command 1 is blank, here is the expected result from command 2.

```
default. supported_cipher_suites = TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
```

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 "/^[^#]*inter_cluster.supported_cipher_suites/
c\inter_cluster.supported_cipher_suites =
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256" /
storage/vcops/user/conf/ssl/secure-communications.properties
```

If the result of command 1 is blank, use the following command to set cipher suites:

```
sed -i "/^[^#]*default.supported_cipher_suites/ c\default.supported_cipher_suites
= TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256" /
storage/vcops/user/conf/ssl/secure-communications.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:

584208	44	-rwsr-xr-x	1	root	root	44696	Feb	4	2019	/usr/bin/su
584210	60	-rwsr-xr-x	1	root	root	54112	Feb	4	2019	/usr/bin/chfn
584646	56	-rwsr-x---	1	root	root	51872	Feb	4	2019	/usr/bin/crontab
584216	40	-rwsr-xr-x	1	root	root	37128	Feb	4	2019	/usr/bin/newgidmap
584206	68	-rwsr-xr-x	1	root	root	63736	Feb	4	2019	/usr/bin/passwd
584211	44	-rwsr-xr-x	1	root	root	44544	Feb	4	2019	/usr/bin/chsh
584218	40	-rwsr-xr-x	1	root	root	37128	Feb	4	2019	/usr/bin/newuidmap
587446	144	-rwsr-xr-x	1	root	root	140856	Feb	4	2019	/usr/bin/sudo
585233	36	-rwsr-xr-x	1	root	root	36144	Feb	4	2019	/usr/bin/umount
584212	32	-rwsr-xr-x	1	root	root	31048	Feb	4	2019	/usr/bin/expiry
584209	76	-rwsr-xr-x	1	root	root	71848	Feb	4	2019	/usr/bin/chage
585231	56	-rwsr-xr-x	1	root	root	52968	Feb	4	2019	/usr/bin/mount
583901	36	-rwsr-xr-x	1	root	root	34944	Feb	4	2019	/usr/bin/
fusermount										
586675	36	-rwsr-xr-x	1	root	root	34952	Feb	4	2019	/usr/bin/
fusermount3										
584217	44	-rwsr-xr-x	1	root	root	44472	Feb	4	2019	/usr/bin/newgrp
584214	80	-rwsr-xr-x	1	root	root	75776	Feb	4	2019	/usr/bin/gpasswd
582975	428	-rwsr-xr-x	1	root	root	432512	Mar	6	2019	/usr/libexec/ssh-
keysign										
587407	80	-rwsr-x---	1	root	root	76224	Feb	4	2019	/usr/libexec/dbus-
daemon-launch-helper										
587109	16	-rwsr-xr-x	1	root	root	14408	Feb	4	2019	/usr/sbin/
usernetctl										
587105	16	-rwxr-sr-x	1	root	root	14384	Feb	4	2019	/usr/sbin/
netreport										
582750	40	-rwsr-xr-x	1	root	root	38960	Feb	4	2019	/usr/sbin/
unix_chkpw										

- 2 Run the `find / -path */proc -prune -o -nouser -print -o -nogroup -print` 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 -not -path "*/sys*" -not -path "*/proc*" -not -path "*/dev*" -perm -o+w | xargs ls -lb` command to verify that none of the files are world writable files by reviewing permissions of all the files on the vApp.

Others should not have write permission. The permissions on these files should be `##4` or `##5`, where `#` equals the default given set of permissions for the Owner and Group, such as `6` or `7`.

- 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.

Apache Configuration

Disable Web Directory Browsing

As a security best practice, ensure that a user cannot browse 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/httpd/httpd.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.

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/httpd/conf/extra/httpd-default.conf | grep ServerTokens` command.

- 2 To modify `ServerTokens Full` to `ServerTokens Prod`, run the `sed -i 's/(ServerTokens\s\+)\sFull/\1Prod/g' /etc/httpd/conf/extra/httpd-default.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 the 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the `install usb-storage /bin/false` 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the line `install bluetooth /bin/false` 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 can 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the following line appears in this file.

```
install sctp /bin/false
```

- 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the DCCP lines appear in the file.

```
install dccp /bin/false
install dccp_ipv4 /bin/false
install dccp_ipv6 /bin/false
```

- 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the `install rds /bin/false` 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the `install tipc /bin/false` 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 can 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the line `install ipx /bin/false` 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 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the line `install appletalk /bin/false` 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 can 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.d/modprobe.conf` file in a text editor.
- 2 Ensure that the line `install decnet /bin/false` 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 necessary.

Procedure

- 1 Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
- 2 Ensure that the line `install ieee1394 /bin/false` 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 2-10. Linux Files and Permissions

Directory or File	Permissions	Groups or Users	Read	Write	Execute
<i>agent directory/bin</i>	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
<i>agent directory/conf</i>	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
<i>agent directory/log</i>	700	Owner	Yes	Yes	No
		Group	No	No	No
		All	No	No	No
<i>agent directory/data</i>	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
<i>agent directory/bin/ep-agent.bat</i>	600	Owner	Yes	Yes	No

Table 2-10. Linux Files and Permissions (continued)

Directory or File	Permissions	Groups or Users	Read	Write	Execute
		Group	No	No	No
		All	No	No	No
<i>agent directory</i> /bin/ep-agent.sh	700	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
<i>agent directory</i> /conf/* (all files in the <code>conf</code> directory)	600	Owner	Yes	Yes	Yes
		Group	No	No	No
		All	No	No	No
<i>agent directory</i> /log/* (all files in the <code>log</code> directory)	600	Owner	Yes	Yes	No
		Group	No	No	No
		All	No	No	No
<i>agent directory</i> /data/* (all files in the <code>data</code> directory)	600	Owner	Yes	Yes	No
		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 2-11. 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		-	-	-	-

Table 2-11. Windows Files and Permissions (continued)

Directory or File	Groups or Users	Full Control	Modify	Read and Execute	Read	Write
<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	-	-	-	-
	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		-	-	-	-

Table 2-11. Windows Files and Permissions (continued)

Directory or File	Groups or Users	Full Control	Modify	Read and Execute	Read	Write
<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 2-12. 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** page 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/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 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 before running the `ep-agent.sh setup` command, 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

Block unnecessary ports on your host server that are not required.

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.

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.
 - d Run `# sysctl -p` to apply the configuration.

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.
 - d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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 can 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.
- d Run `# sysctl -p` to apply the configuration.

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 | grep "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 `net.ipv4.conf.default.accept_source_route=0` are set to 0.
 - c Save and close the file.
 - d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

Configure the Host System to Use IPv4 TCP SYN Cookies

As a security best practice, verify that the host system uses IPv4 Transmission Control Protocol (TCP) SYN cookies. 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. SYN cookies 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 defense 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 SYN cookies.

2 Configure the host system to use IPv4 TCP SYN cookies.

- 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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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 accepts 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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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 that 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.
- d Run `# sysctl -p` to apply the configuration.

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.
- d Run `# sysctl -p` to apply the configuration.

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. The ports should be allowed/opened in local network for vRealize Operations Manager inter-node communication and for customer to vRealize Operations Manager communication.

Table 2-13. 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 primary node.
5433	TCP	Used by the primary 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.

Table 2-13. Minimum Required Incoming Ports (continued)

Port	Protocol	Comments
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.
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 2-14. 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.

Cipher Suites and Protocols

The cipher suites and relevant protocols are listed when FIPS is in On and Off mode.

Cipher Suites When FIPS is On

Here are the cipher suites lists when FIPS is On. The cipher suites are classified based on incoming, internode, and outbound connections. The cipher suite list is a comma-separated list.

Incoming Connections to vRealize Operations Manager

Table 2-15. Cipher Suites for Incoming Connections

Name	Cipher Suites
Configured Cipher Suites	
Apache Ciphers Protocol - TLS 1.2	ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA, ECDHE-RSA-AES128-SHA, AES256-GCM-SHA384, AES128-GCM-SHA256, AES256-SHA256, AES128-SHA256, AES256-SHA, AES128-SHA
What you can configure: To find Apache relays to the OS cipher suite list, run the CLI command: <code>openssl ciphers -v</code>	

Internode Connections between vRealize Operations Manager Nodes

Table 2-16. Cipher Suites for Internode Connections

Name	Cipher Suites
Configured Cipher Suites	
inter_cluster Protocol - TLSv1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
What you can configure:	

Table 2-16. Cipher Suites for Internode Connections (continued)

Name	Cipher Suites
All the possible cipher suites for internode connections.	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA
Note The PostgreSQL and Cassandra cipher suite lists must have an intersection with the inter_node cipher suite list. The inter_node proper cipher suite selection will avoid PostgreSQL and Cassandra from non-secure cipher suite usage.	

Outbound Connections from vRealize Operations Manager

Outbound cipher suites that are configured are classified into three types:

- Adapter to Source
- Authentication Sources
- Outbound Plugins

Table 2-17. Adapter to Source

Name	Cipher Suites
All Adapters Protocols - TLSv1.2, TLSv1.1, TLSv1	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-18. Authentication Sources

Name	Cipher Suites
vIDM Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA
sso_util Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-18. Authentication Sources (continued)

Name	Cipher Suites
csp Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA
LDAP Protocol - TLSv1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-19. Outbound Plugins

Name	Cipher Suites
cprc_connection Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA
marketplace_manager Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-19. Outbound Plugins (continued)

Name	Cipher Suites
email_sender Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-19. Outbound Plugins (continued)

Name	Cipher Suites
rest_sender Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA
lint_rest_template Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Table 2-20. Outbound Cipher Suites that You Can Configure

Name	Cipher Suites
All the possible cipher suites you can configure for an outbound connection.	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA

Cipher Suites When FIPS is Off

Here are the lists of cipher suites when FIPS is Off. The cipher suites are classified based on incoming, internode, and outbound connections. The cipher suite list is a comma-separated list.

Incoming Connections to vRealize Operations Manager

Table 2-21. Cipher Suites for Incoming Connections

Name	Cipher Suites
Configured Cipher Suites	
Apache Ciphers Protocol - TLS 1.2	ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA, ECDHE-RSA-AES128-SHA, AES256-GCM-SHA384, AES128-GCM-SHA256, AES256-SHA256, AES128-SHA256, AES256-SHA, AES128-SHA
What you can configure: To find Apache relays to the OS cipher suite list, run the CLI command: <code>openssl ciphers -v</code> .	

Internode Connections between vRealize Operations Manager Nodes

Table 2-22. Cipher Suites for Internode Connections

Name	Cipher Suites
Configured Cipher Suites	
inter_cluster Protocol - TLSv1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
What you can configure:	

Table 2-22. Cipher Suites for Internode Connections (continued)

Name	Cipher Suites
All the possible cipher suites for internode connections.	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV
Note The PostgreSQL and Cassandra cipher suite lists must have an intersection with the inter_node cipher suite list. The inter_node proper cipher suite selection will avoid PostgreSQL and Cassandra from non-secure cipher suite usage.	

Outbound Connections from vRealize Operations Manager

Outbound cipher suites that are configured are classified into three types:

- Adapter to Source

- Authentication Sources
- Outbound Plugins

Table 2-23. Adapter to Source

Name	Cipher Suites
All adapters Protocols - TLSv1.2, TLSv1.1, TLSv1	TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_DH_RSA_WITH_AES_256_CBC_SHA, TLS_DH_DSS_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_DH_RSA_WITH_AES_128_CBC_SHA, TLS_DH_DSS_WITH_AES_128_CBC_SHA

Table 2-24. Authentication Sources

Name	Cipher Suites
vIDM Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV
sso_util Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,

Table 2-24. Authentication Sources (continued)

Name	Cipher Suites
	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV

Table 2-24. Authentication Sources (continued)

Name	Cipher Suites
csp Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV
LDAP Protocol - TLSv1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,

Table 2-24. Authentication Sources (continued)

Name	Cipher Suites
	TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_DH_RSA_WITH_AES_128_CBC_SHA, TLS_DH_DSS_WITH_AES_128_CBC_SHA, TLS_DH_RSA_WITH_AES_256_CBC_SHA, TLS_DH_DSS_WITH_AES_256_CBC_SHA

Table 2-25. Outbound Plugins

Name	Cipher Suites
cprc_connection Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV
marketplace_manager Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,

Table 2-25. Outbound Plugins (continued)

Name	Cipher Suites
	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV

Table 2-25. Outbound Plugins (continued)

Name	Cipher Suites
email_sender Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV

Table 2-25. Outbound Plugins (continued)

Name	Cipher Suites
rest_sender Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV
lint_rest_template Protocol - TLSv1.2	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,

Table 2-25. Outbound Plugins (continued)

Name	Cipher Suites
	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV

Table 2-26. Outbound Cipher Suites that You Can Configure

Name	Cipher Suites
All the possible cipher suites you can configure for an outbound connection.	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_EMPTY_RENEGOTIATION_INFO_SCSV

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

Install VMware vRealize Operations Manager to create and configure one or more nodes that collect and analyze object data from your environment.

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](#)
- [Upgrade, Backup and Restore](#)

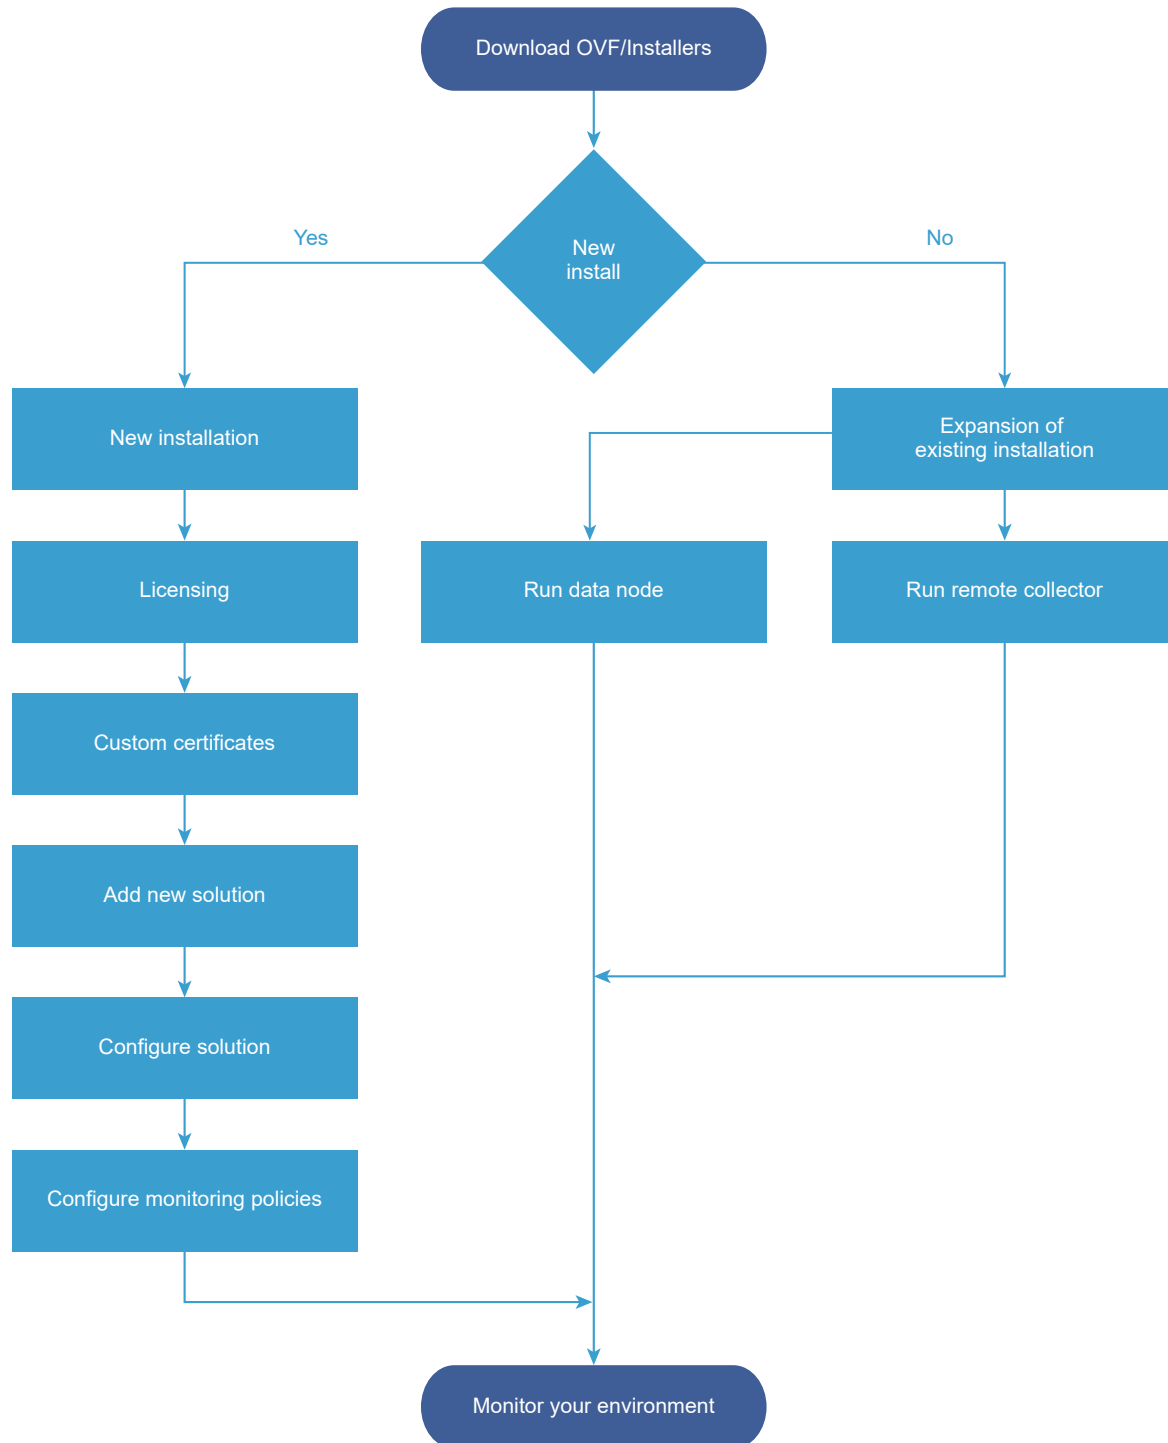
About Installing

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, 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



To automate installation, configuration, upgrade, patch, configuration management, drift remediation and health from within a single pane of glass, you can use vRealize Suite Lifecycle Manager. If you are a new user, click [here](#) to install **vRealize Suite Lifecycle Manager**. This provides the IT Managers of Cloud admin resources to focus on business-critical initiatives, while improving time to value (TTV), reliability, and consistency.

You can also install upgrade vRealize Operations Manager by using vRealize Suite Lifecycle Manager. For more information, see the [Creating an Environment from Configure vRealize Products](#).

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 <i>solution</i> , <i>management pack</i> , <i>adapter</i> , and <i>plug-in</i> 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 (primary 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 extends larger.

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

Master Node

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

In a single-node installation, the primary 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 primary 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 primary node.

The following node types are 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 the 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 might be installed on a different operating system than the rest of the cluster.

Witness Node

To use vRealize Operations Manager continuous availability (CA), the cluster requires that you have a witness node. If the network connection between the two fault domains is lost, the witness node acts as a decision maker regarding the availability of vRealize Operations Manager.

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 purposes. Unlike the data nodes, the remote collector nodes only perform the collector role of vRealize Operations Manager. These remote collectors do not store data or process any analytics functions. Remote collectors collect data from integrated objects and then forward the data back to the primary node. The primary node then processes the data which you then view as reports and analytics.

Remote collectors are very useful when you have multiple locations. You can deploy remote collectors on remote location sites and only deploy the primary node at the primary location.

You must have at least one primary node before adding remote collector nodes.

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. To deploy a remote collector node, see [Run the Setup Wizard to Create a Remote Collector Node](#).

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.

Ports information for vRealize Operations Manager is available on [Ports and Protocol](#).

About vRealize Operations Manager High Availability

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

With HA, data stored in the primary 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 primary node. If you have more than one data node, the data stored in the primary node can be stored and replicated in any of the other nodes. But in case the primary node fails, only the replica node can function as the replacement of the primary 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 primary provides, were the primary to fail for any reason. If the primary 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 primary node problem causes failover, the replica node becomes the primary 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 primary node. When a primary node exits an HA-enabled cluster, primary 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.
- Remove the failed primary node then re-enable HA by converting a data node into replica. Removed primary nodes cannot be repaired and readded to vRealize Operations Manager.
- Remove the old, failed primary node and then change to non-HA operation by disabling HA. Removed primary nodes cannot be repaired and readded to vRealize Operations Manager.

- In the administration interface, after an HA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary 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 primary/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, and the replica backup of the primary 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 might cause the loss of more than one node, which is not supported, and all of vRealize Operations Manager can become unavailable.

The opposite is also true. Without HA, you can keep nodes on the same host, and it will not make a difference. Without HA, the loss of even one node can 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 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 displays a different quantity of objects from that shown in the Administration UI.

About vRealize Operations Manager Continuous Availability

vRealize Operations Manager supports continuous availability (CA). CA separates the vRealize Operations Manager cluster into two fault domains, stretching across vSphere clusters, and protects the analytics cluster against the loss of an entire fault domain.

You can configure the analytics cluster with Continuous Availability. This allows the cluster nodes to be stretched across two fault-domains. A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. With CA, the two fault domains permit vRealize Operations Manager to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

To enable continuous availability within vRealize Operations Manager, the witness node must be deployed in the cluster. The witness node does not collect nor store data. In a situation where network connectivity between the two fault-domains is lost, the cluster would go into a split-brain situation. This situation is detected by the Witness Node and one of the fault domains will go offline to avoid data inconsistency issues. You will see a **Bring Online** button on the admin UI of the nodes which are made offline by the witness node. Before using this option to bring the fault domain online, ensure that the network connectivity between the nodes across the two fault domains is restored and stable. Once confirmed you can bring the fault domain online.

With CA, the data stored in the primary node and data nodes grouped in fault domain 1 is always 100% synced to the replica node and data nodes paired in fault domain 2. To enable CA, you must have at least one data node deployed, in addition to the primary node. If you have more than one data node, there must be an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes based on the appropriate sizing requirements. The data stored in the primary node in fault domain 1 is stored and replicated in the replica node in fault domain 2. The data stored in the data nodes in fault domain 1 is stored and replicated in the paired data nodes in fault domain 2. But in case the primary node fails, only the replica node can function as the replacement of the primary node.

- CA protects the analytics cluster against the loss of half the analytics nodes specific to one fault domain. You can stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When CA is enabled, the replica node can take over all functions that the primary node provides, in case of a primary node failure. The 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.

Note In case of a primary node failure, the replica node becomes the primary node, and the cluster runs in degraded mode. To fix this, perform any one of the following actions.

- Correct the primary node failure manually.
 - Return to CA mode by replacing the primary node. Replacement nodes do not repair the node failure, instead a new node assumes the primary node role.
-
- In the administration interface, after a CA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary node. To refresh the display and enable the removal of the node, refresh the browser.

- When CA is enabled, the cluster can survive the loss of half the data nodes, all in one fault domain, without losing any data. CA protects against the loss of only one fault domain at a time. Simultaneously losing data and primary/replica nodes, or two or more data nodes in both fault domains, is not supported.
- A CA enabled cluster will be non-functional if you power off the primary node or the primary node replica while one of the fault domains is down.
- When CA is enabled, it lowers the vRealize Operations Manager capacity and processing by half, because CA creates a redundant copy of data throughout the cluster, and the replica backup of the primary node. Consider your potential use of CA when planning the number and size of your vRealize Operations Manager cluster nodes. See [Sizing the vRealize Operations Manager Cluster](#).
- When CA is enabled, deploy analytics cluster nodes, in each fault domain, on separate hosts for redundancy and isolation. You can also use anti-affinity rules that keep nodes on specific hosts in the vSphere clusters.
- If you cannot keep the nodes separate in each fault domain, you can still enable CA. A host fault might cause the loss of the data nodes in the fault domain, and vRealize Operations Manager can still be available in the other fault domain.
- If you cannot split the data nodes into different vSphere clusters, do not enable CA. A cluster failure can cause the loss of more than half of the data nodes, which is not supported, and all of vSphere might become unavailable.
- Without CA, you can keep nodes on the same host in the same vSphere. Without CA, the loss of even one node might make all of vRealize Operations Manager unavailable.
- When you power off data nodes in both fault domains and change the network settings of the VMs, it affects the IP address of the data nodes. After this point, the CA cluster is no longer accessible and all the nodes status change to "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 CA-enabled cluster, one or more vCenter adapters associated with that node stops collecting. You must change the adapter configuration to pin them to another node before removing the node.
- The administration interface displays the resource cache count, which is created for active objects only, but the inventory displays all objects. When you remove a node from a CA-enabled cluster allowing the vCenter adapters to collect data and rebalance each node, the inventory displays a different quantity of objects from that shown in the administration interface.

Preparing for Installation

When you prepare for your installation, consider some of these best practices, cluster, sizing and scaling 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 both, Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). You can use IPv4 or IPv6 or both. If the environment has a dual-stack support with both IPv4 and IPv6 protocols, all nodes in the cluster must follow the same protocol. When using IPv6, the **Prefer IPv6** flag must be enabled during the OVF deployment for each node. If you set the **Prefer IPv6** flag, vRealize Operations Manager uses IPv6 for its internal communications. It does not affect how vRealize Operations Manager handles its external communications. The use of IPv6 with vRealize Operations Manager requires that certain limitations be observed.

Considerations While Using IPv6

- All vRealize Operations Manager cluster nodes, including remote collectors, must have IPv6 addresses. Do not mix IPv6 and IPv4.
- 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. Primary nodes and replica nodes still require fixed addresses, which are true for IPv4 as well.
- Your DNS server must be configured to support IPv6.
- When adding nodes to the cluster, enter the IPv6 address of the primary 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 When vRealize Operations Manager is using IPv6, vCenter Server might 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**
- **Witness Node Deployment Type.** The witness node must be the same vApp deployment.
- **Analytics Cluster Node Sizing.** In the analytics cluster, CPU, memory, and disk size must be identical for all nodes.

Primary, 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.
- **Witness Node Sizing.** The witness node has only one size and may be of different sizes from remote collectors 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.

- **Witness Node Placement.** You may place the witness node in a different vSphere cluster separate from the analytics nodes.
- **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](#) .
- If you expect to enable CA, place analytics cluster nodes on separate hosts in fault domains, stretched across vSphere clusters. See [About vRealize Operations Manager Continuous 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 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 the 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 primary and replica nodes must use a 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.
- With Continuous Availability enabled, separate analytics cluster nodes into fault domains, stretched across vSphere clusters
- Packet Round Trip Time 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 Primary, 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 primary 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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 1 From the Configuration Parameters, remove the setting <code>numa.vcpu.preferHT</code> 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) or CA (Continuous 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.

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 primary node, reserve a static IP address for the virtual machine, and know the associated domain name, domain search path, domain name servers, 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/CA replica node, reserve a static IP address for the virtual machine, and store the associated domain name, domain search path, domain name servers, default gateway, and network mask values for later use.

In addition, familiarize yourself with HA node placement as described in [About vRealize Operations Manager High Availability](#) and CA node allocation as described in [About vRealize Operations Manager Continuous Availability](#).

- Plan your domain and machine naming so that the deployed virtual machine name begins and ends with an 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.

- Plan 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** **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 the 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, see [Add Data Disk Space to a vRealize Operations Manager vApp Node](#).

- 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 Under Networking Properties, in case of a static IP, specify the associated **Default Gateway**, **Domain Name**, **Domain Search Path**, **Domain Name Servers**, **Network 1 IP Address**, and **Network 1 Netmask** values. In case of DHCP, leave all the fields blank. The primary node and replica node require a static IP. A data node or remote collector node can use DHCP or a static IP.

Note The hostname is configured using DHCP and DNS. If a static IP is used the hostname is configured according to the node name specified during node configuration, after deployment.

- 11 In the 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.

Note You cannot configure nodes to different time zones.

- 12 (Optional) In Properties, under Application, select the option for IPv6 .
- 13 (Optional) If you want to deploy a FIPS enabled vRealize Operations Manager setup, in the FIPS setting, select the **Enable FIPS Mode** check box.
- 14 Click **Next**.
- 15 Review the settings and click **Finish**.
- 16 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 primary node, a data node, a high availability primary replica node, or a remote collector node. The primary 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 page



Perform a New Installation on the vRealize Operations Manager Product UI

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

For example: **Ops-Master**

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

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 primary node and replica node.

10 Click **Next**.

11 Configure the vRealize Operations Manager availability. To install vRealize Operations Manager with availability, enable the **Availability Mode** and select High Availability or Continuous Availability. To continue your installation on full capacity, click **Next**.

Note You can enable High Availability or Continuous Availability after installation from the administrator interface.

12 Click the Add icon to add a node.

a Enter the **Node Name** and **Node Address**.

b Select the **Current Cluster Role**.

Note This step is optional if you use the default configuration. If you select High Availability for this cluster option, you can select a node from the added list of nodes to be the replica node. Although, only one node from the list can be selected as a replica node. For more information on High Availability, see [Adding High Availability to vRealize Operations Manager](#). If you select Continuous Availability for this cluster, add at least one witness node and an even number of data nodes including the primary node and divide them across two fault domains. For more information, see [Adding Continuous Availability](#).

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

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

Results

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

What to do next

After creating the primary 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 primary node that is the required, initial node in your vRealize Operations Manager cluster.

The primary node performs administration for the cluster and must be online before you configure any new nodes. In addition, the primary node must be online before other nodes are brought online. If the primary node and replica node go offline together, bring them back online separately. Bring the primary node online first, and then bring the replica node online.

Advantages of a New Installation

You can use the new installation to create a primary node during the first installation of vRealize Operations Manager. With the primary 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 is on the same primary node. A multiple-node cluster includes one primary 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 continuous availability, you need a witness node and an even number of data nodes including the primary node. For more information on creating a primary 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 primary 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 a new installation. Do 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 primary node. Select express installation option when installing for the first time.

Prerequisites

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

Procedure

- 1 Navigate to the name or IP address of the node that will be the primary 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 primary 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 primary node. The express installation uses the default certificates, which differ 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 primary 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 primary 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 primary 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 primary node.
- Note the fully qualified domain name (FQDN) or IP address of the primary 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 primary 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 primary replica node.
 - Navigate to the vRealize Operations Manager administrator interface and log in <https://vROps-IP/admin>. Verify that all the nodes are listed under the **Nodes in the vRealize Operations Manager Cluster**. Then, click **Start vRealize Operations Manager** to start the cluster and 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 primary 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 primary node before you add data nodes to form a cluster.

Installing vRealize Operations Manager on VMware Cloud on AWS

You can use your on-premises vRealize Operations Manager to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into

vRealize Operations Manager. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations Manager on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations Manager on-premises and VMware Cloud.
- Scale the existing vRealize Operations Manager cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- The `cloudadmin@vmc.local` user in VMware Cloud has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines on VMware Cloud. Active and consumed memory utilizations continue to work in this case.
- Cost computation is disabled on the VMware Cloud on AWS inventory since the cost model is different from the on-premises vCenter Server. The cost of infrastructure coming from VMware Cloud on AWS can be managed using vRealize Operations Manager management pack for VMware Cloud on AWS.
- The compliance workflows in vRealize Operations Manager 8.3 work for virtual machines running on a vCenter Server in VMware Cloud on AWS. The compliance checks for VMware management objects such as Hosts, vCenter, and so on, are not available.
- Only migration planning scenarios with VMware Cloud are supported.
- Workload optimization including pDRS and host-based business intent does not work because VMware managers cluster configurations.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with vRealize Operations Manager 8.3. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud does not support vRealize Operations Manager plugin.
- You cannot log in to vRealize Operations Manager using your VMware Cloud vCenter Server credentials.
- Only migration planning scenarios with VMware Cloud are supported.

Using vRealize Operations Manager on-premises on VMware Cloud on AWS

Extend the monitoring capabilities of your on-premises vRealize Operations Manager to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations Manager. Create an adapter instance both for vCenter Server and VMware vSAN to collect data from VMware Cloud and bring that into vRealize Operations Manager. You can either connect directly to the vCenter Server or use a remote collector which

can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations Manager primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations Manager remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Datacenter level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations Manager primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations Manager cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note In case of a vCenter adapter instance, set the **Cloud Type** to **VMware Cloud on AWS**.

Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-6. vRealize Operations On-Premises collecting data from VMware Cloud and AWS without remote data collectors





Figure 3-7. vRealize Operations On-Premises collecting data from VMware Cloud and AWS with remote data collectors



Deploying vRealize Operations Manager on VMware Cloud on AWS

If you have moved a large part of your environment into VMware Cloud, you can deploy or migrate your vRealize Operations Manager instance into VMware Cloud directly. After the vRealize Operations Manager cluster is deployed on VMware Cloud, you can collect data from other VMware Cloud SDDCs and the SDDC located on-prem using remote collectors. You can deploy remote collectors to send over data into the centralized analytics cluster deployed in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations Manager cluster in VMware Cloud, see [Deployment of vRealize Operations Manager](#).

Note Deploy the OVF template in the VMware Cloud on the data center level. VMware Cloud has two resource pools, the regular workload and the administrative workload. You can only deploy the new OVF template in the workload resource pool.

- 2 Deploy the remote collectors in vRealize Operations Manager , see [Create a Remote Collector](#).

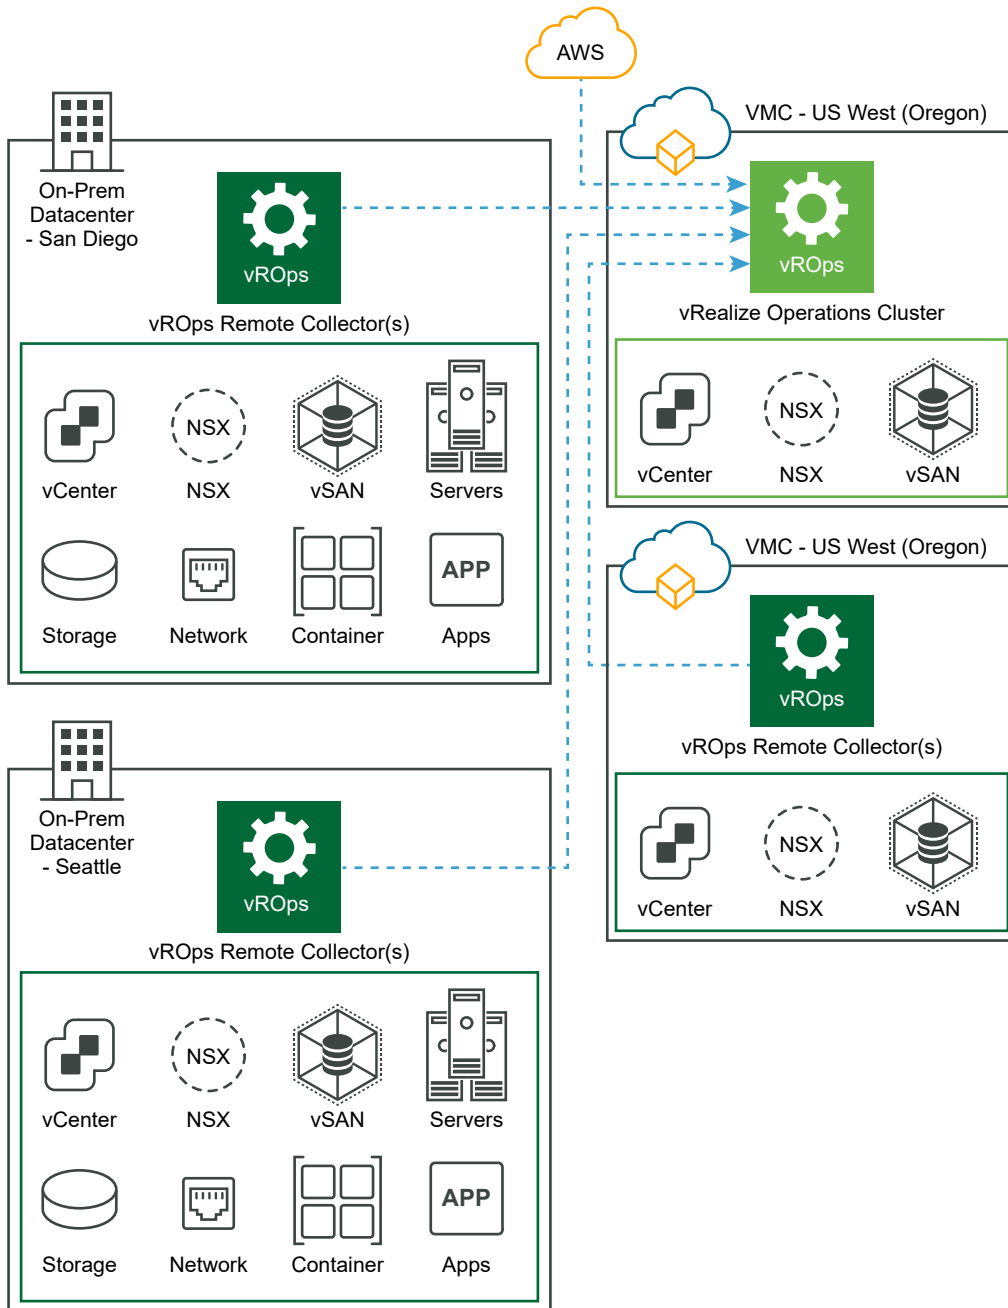
Note VMware Cloud is set in an isolated network and so, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations Manager primary node and the remote collector you have created. To do so, you can use a VPN or a direct connection with no NAT.

- 3 Add and configure an adapter instance in the vRealize Operations Manager cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note If the Remote collectors are deployed on-premises, set **Cloud Type** to **Private Cloud**. However, if you deploy remote collectors in another VMware Cloud, set the **Cloud Type** to **VMware Cloud on AWS**.

Ensure that the remote collector is assigned to the adapter instance and the data collection of the adapter instance happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-8. vRealize Operations in VMware Cloud collecting data from other VMware Cloud SDDC, AWS and On-Premise with remote data collectors



Installing vRealize Operations Manager for Azure VMware Solution

You can use your on-premises vRealize Operations Manager to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations Manager. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations Manager on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations Manager on-premises and VMware Cloud.
- Scale the existing vRealize Operations Manager cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Microsoft manages the compliance of Azure VMware Solution hosts. Ignore the compliance alerts for Azure VMware Solution hosts.
- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters might appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation is not supported on Azure VMware Solution. Ignore all the cost metrics.
- The end-user on the vCenter Server on Azure VMware Solution has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations Manager using the credentials of the vCenter Server on Azure VMware Solution.
- The vCenter Server on Azure VMware Solution does not support the vRealize Operations Manager plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.
- Management VMs are hidden from the end user visibility while the respective VMDKs are not. As a result, vRealize Operations Manager considers management VMDKs as orphaned and should be ignored.

Using vRealize Operations Manager on-premises for Azure VMware Solution

Extend the monitoring capabilities of your on-premises vRealize Operations Manager to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations Manager. Create an adapter instance both for vCenter Server and VMware vSAN to collect data from VMware Cloud and bring that into vRealize Operations Manager. You can either connect directly to the vCenter Server or use a remote collector which can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations Manager primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations Manager remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations Manager primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations Manager cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-9. (Recommended) vRealize Operations On-Premises collecting data from Azure VMware Solution with remote data collectors





Figure 3-10. vRealize Operations On-Premises collecting data from Azure VMware Solution without remote data collectors



Deploying vRealize Operations Manager on Azure VMware Solution

Deployment of vRealize Operations Manager on Azure VMware Solution is not supported.

Installing vRealize Operations Manager for Google Cloud VMware Engine

You can use your on-premises vRealize Operations Manager to manage and monitor your cloud infrastructure on VMware Cloud by simply adding your VMware Cloud based vCenter Server into vRealize Operations Manager. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of vRealize Operations Manager on to VMware Cloud. It provides you with a hybrid view of your environment.

Prerequisites

- A VPN or a direct connection to set up the bidirectional access between the nodes and remote collectors of vRealize Operations Manager on-premises and VMware Cloud.
- Scale the existing vRealize Operations Manager cluster before adding the new VMware Cloud SDDC sites. To get the appropriate sizing, see [vRealize Operations Manager Online Sizer](#).

Known Limitations

- Google manages the compliance of Google Cloud VMware Engine hosts. Ignore the compliance alerts for Google Cloud VMware Engine hosts.

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters may appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation is not supported on Google Cloud VMware Engine. Ignore all the cost metrics.
- The end-user on the vCenter Server on Google Cloud VMware Engine has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations Manager using the credentials of the vCenter Server on Google Cloud VMware Engine.
- The vCenter Server on Google Cloud VMware Engine does not support the vRealize Operations Manager plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.
- Management VMs are hidden from end user visibility while the respective VMDKs are not. As a result, vRealize Operations Manager considers management VMDKs as orphaned and should be ignored.

Using vRealize Operations Manager on-premises for Google Cloud VMware Engine

Extend the monitoring capabilities of your on-premises vRealize Operations Manager to monitor the VMware Cloud vCenter Server by connecting the VMware Cloud vCenter Server as an end point inside vRealize Operations Manager. Create an adapter instance both for vCenter Server and VMware vSAN to collect data from VMware Cloud and bring that into vRealize Operations Manager. You can either connect directly to the vCenter Server or use a remote collector which can be deployed inside a VMware Cloud SDDC to ensure that the data can be compressed and encrypted.

Note If the network latency between vRealize Operations Manager primary node and VMware Cloud is greater than 5 milliseconds, you should deploy remote collectors in VMware Cloud.

Procedure

- 1 Deploy the vRealize Operations Manager remote collectors in VMware Cloud, see [Create a Remote Collector](#).

Note Deploy the OVF in the SDDC-Data Center level and select the **Compute Resource Pools** and validate your deployment. You can only select the workload datastore for storage when deploying the OVF in VMware Cloud.

Since VMware Cloud is set in an isolated network, the remote collectors cannot view or connect to the primary node. To collect data, you must set up the bidirectional access between the vRealize Operations Manager primary node and the remote collectors you have created. To do so, you can use a VPN or create a direct connection with no-NAT.

- 2 Add and configure an adapter instance in the vRealize Operations Manager cluster in VMware Cloud. To configure a vCenter adapter, see [Configure a vCenter Adapter Instance in vRealize Operations Manager](#). To configure a vSAN adapter, see [Configure a vSAN Adapter Instance](#).

Note Ensure that the remote collector is assigned to the adapter instance and the data collection happens through the remote collectors that you have set up. Select the newly deployed remote collectors for **Collectors/Groups** under **Advanced Settings**.

Figure 3-11. (Recommended) vRealize Operations On-Premises collecting data from Google Cloud VMware Engine with remote data


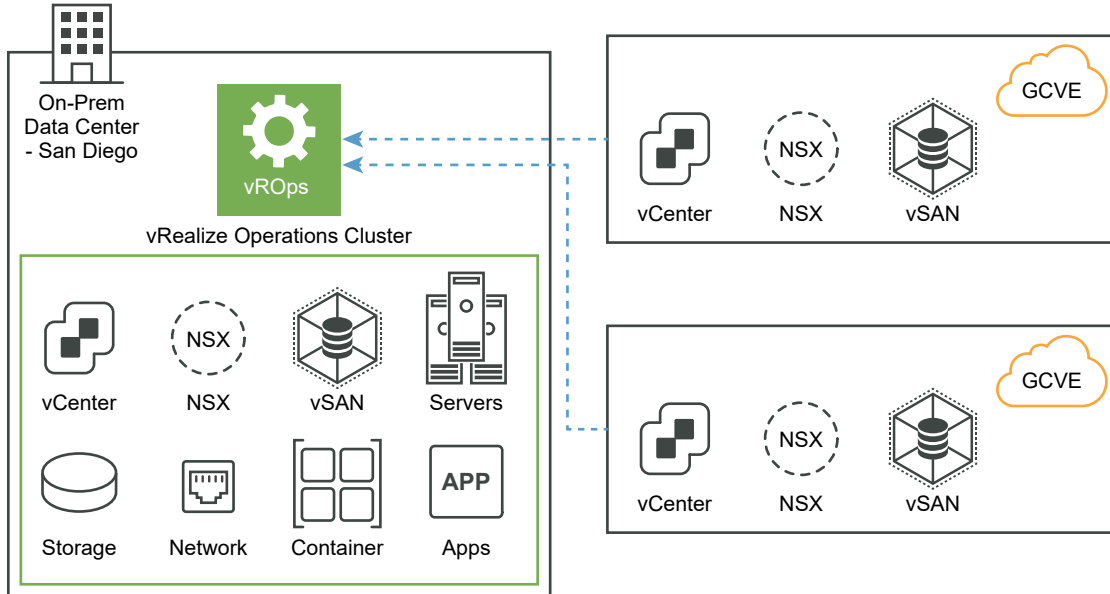


Figure 3-12. vRealize Operations On-Premises collecting data from Google Cloud VMware Engine without remote data collectors



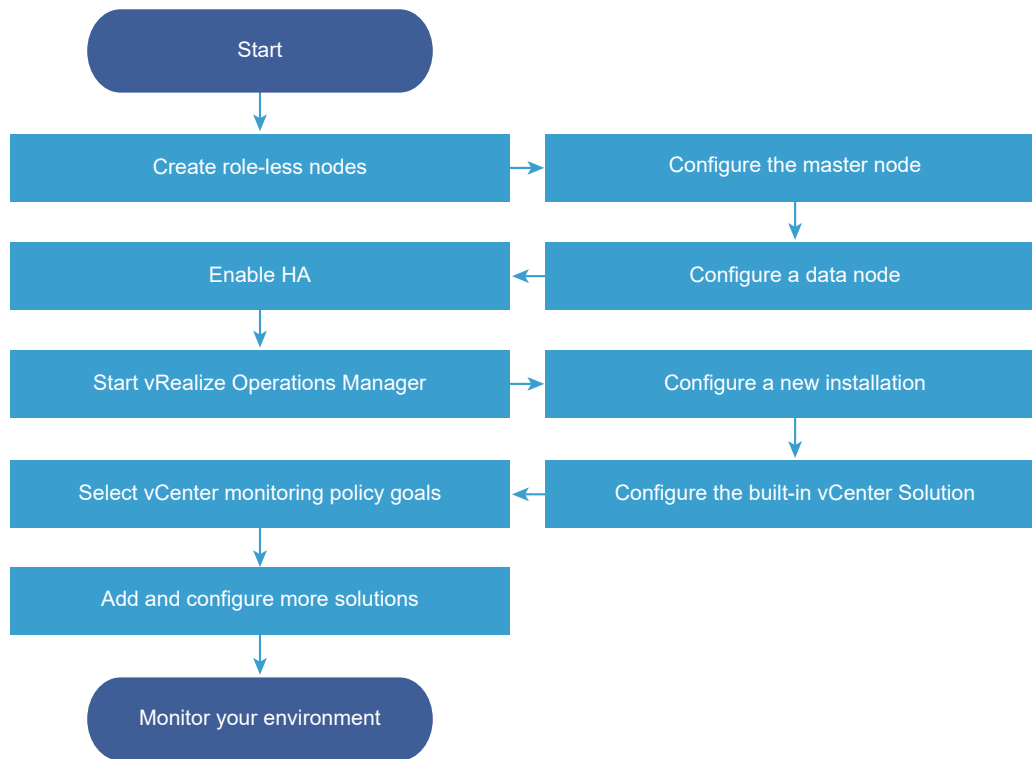
Deploying vRealize Operations Manager on Google Cloud VMware Engine

Deployment of vRealize Operations Manager on Google Cloud VMware Engine is not supported.

Resize your Cluster by Adding Nodes

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

Figure 3-13. 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 primary node.
- Note the fully qualified domain name (FQDN) or an IP address of the primary 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 primary 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 primary 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 primary 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 primary node.

Run the Setup Wizard to Add a Primary Replica Node

To enable high availability (HA) for a vRealize Operations Manager cluster, specify one of the data nodes to become a replica of the primary node.

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

You can 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 primary node.
- Create and configure a data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary 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 user name 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 primary 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 primary node and replica node go offline, and the primary remains offline for any reason while the replica goes online, the replica node does not take over the primary 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 primary node and bring the remaining cluster nodes online.

What to do next

After creating a primary 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.

Adding Continuous Availability

Continuous availability prevents data loss in the event of one or more node failures. This mode requires one witness node, one primary node, and one data node divided across two fault domains. The witness node lies outside the fault domains. By default, the primary node is assigned to **Fault Domain 1**. The data node becomes the replica node and is assigned to **Fault Domain 2**. The primary node and the replica node create a pair. The number of data nodes including the primary node should always be an even number not exceeding 16. Each data node added to **Fault Domain 1** must have a pair in **Fault Domain 2** to preserve and replicate data that is added to its peer.

Enable Continuous Availability in vRealize Operations Manager

You can enable continuous availability (CA) for vRealize Operations Manager to protect your data if there is one or more node failures.

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

You can enable CA in the vRealize Operations Manager cluster at the installation time or after vRealize Operations Manager is up and running. Adding CA 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 primary node.
- Create and configure the witness node.

Note While deploying an OVA file, you can select the recommended CPU/RAM configuration for the witness node.

- Create and configure one data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary 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 user name of **admin**.
- 3 Enter the vRealize Operations Manager administrator password and click **Log In**.
- 4 Under Continuous Availability, click **Enable CA**.

The Continuous Availability wizard opens. The Witness node exists outside the fault domains. The primary node is already assigned to **Fault Domain 1**.

Note You can enter names for each Fault Domain during installation. You can also edit the fault domain names after enabling continuous availability.

- 5 To create a pair with the primary node, drag the data nodes to **Fault Domain 2**.

Note You can add a maximum of 16 data nodes including the primary node and divide them between the fault domains to create eight pairs. You can also add remote collector nodes outside the fault domains as required.

- 6 Click **Ok**.

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, fault domains, or individual nodes, enabling or disabling high availability (HA) or continuous availability (CA), reviewing statistics related to the installed adapters, and rebalancing the workload for a 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	<p>You can change the status of a node to online or offline.</p> <p>In a high availability (HA) cluster, taking the primary or replica offline causes vRealize Operations Manager to run from the remaining node and for HA status to be degraded.</p> <p>In continuous availability (CA) cluster, taking the primary or replica offline causes vRealize Operations Manager to run in a degraded status.</p> <hr/> <p>Note You cannot convert a High Availability (HA) enabled cluster to a Continuous Availability cluster and vice versa. You must first disable the cluster availability, so that the cluster becomes a standard cluster and then enable HA or CA as required.</p> <hr/> <p>Any manual or system action that restarts the cluster brings all vRealize Operations Manager nodes online, including any nodes that you had taken offline.</p> <p>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 and click the Start Collector icon .</p>
Enable or Disable High Availability	Administration	<p>Enabling 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.</p> <p>To enable high availability, see Adding High Availability to vRealize Operations Manager.</p> <p>Disabling high availability restarts the vRealize Operations Manager cluster.</p> <p>After you disable high availability, the replica node in vRealize Operations Manager converts back to a data node and restarts the cluster.</p>

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

Procedure	Interface	Description
Enable or Disable Continuous Availability	Administration	<p>Enabling continuous availability requires the cluster to have at least one witness node, and at least two data node, with all nodes online or all offline. You cannot use Remote Collector nodes. To enable continuous availability, see Adding Continuous Availability.</p> <p>Disabling continuous availability restarts the vRealize Operations Manager cluster.</p> <p>When you disable continuous availability, you can choose to keep all your nodes or cut out one of the fault domains.</p> <ul style="list-style-type: none"> Click Simply Disable with keeping all nodes to keep all your nodes when you disable continuous availability. <p>Note You cannot disable continuous availability if one of your nodes is faulty. If you want to keep all your nodes, you must fix or replace the faulty node before you proceed.</p> <ul style="list-style-type: none"> Click Cut-Out one Fault Domain and then select the fault domain you want to keep. The other fault domain and the witness node are deleted. <p>After you disable continuous availability, the replica node in vRealize Operations Manager converts back to a data node and restarts the cluster.</p>
Add Nodes	Administration	<p>You can add one or more nodes for your cluster. In a FIPS enabled environment, new nodes must be FIPS compliant. In a FIPS disabled environment, new nodes must be FIPS disabled. Enabling continuous availability requires one witness node, and an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes.</p>
Replace Nodes	Administration	<p>You can add nodes and replace them with a downed or non-functional node in a cluster.</p>
Generate Passphrase	Administration	<p>You can generate a passphrase to use instead of the administrator credentials to add a node to this cluster.</p> <p>The passphrase is only valid for a single use.</p>

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

Procedure	Interface	Description
Remove a Node	Administration	<p>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.</p> <p>You must not re-add nodes to vRealize Operations Manager that you already removed. If your environment requires more nodes, add new nodes instead.</p> <p>When you perform maintenance and migration procedures, you should take the node offline, not remove the node.</p>
Configure NTP	Product	The nodes in vRealize Operations Manager cluster synchronize with each other by standardizing on the primary 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 includes a central page where you can monitor and manage the nodes in your vRealize Operations cluster and 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 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 cluster. Once CA is enabled, it displays the status of the two fault domains.
High Availability	Indicates whether HA is enabled, disabled, or degraded.
Continuous Availability	Indicates whether CA is enabled, disabled, or degraded.

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

Table 3-5. Nodes in the vRealize Operations Cluster

Option	Description
Node Name	Machine name of the node. The node that you are logged into displays a dot next to the name.
Node Address	Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes can use DHCP or static IP.
Cluster Role	Type of vRealize Operations node: primary, data, replica, or remote collector.
Fault Domain	Displays the fault domain a node is associated to in a CA enabled cluster. Note This column appears only if CA is enabled.
Node Pair	Displays which pair the node belongs to. For example, in CA, nodes are added in pairs. If there are four nodes, the column displays whether the node is part of pair one or two. Note This column appears only if CA is enabled.
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.
Objects Being Collected	Total environment objects that the node collected.
Metrics in Process	Total metrics that the node has discovered since being added to the cluster.

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

Option	Description
Metrics Being Collected	Total metrics the node has collected since being added to the cluster.
Version	Displays the vRealize Operations software version and the build number installed on the node.

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.

Troubleshooting

Troubleshooting Cluster Problems

A multi-node vRealize Operations Manager cluster does not behave as expected.

Problem

A multi-node vRealize Operations Manager cluster does not behave as expected because of general problems within the cluster or because of suspected firewall concerns.

The problems might occur because of multiple reasons:

- You may be unable to install or uninstall management packs.
- The node shows as offline in the user interface even though it is online.
- You might face problems with new nodes joining the cluster.

Solution

Login to each vRealize Operations Manager node in the cluster and run the following script:

```
$VMWARE_PYTHON_3_BIN /usr/lib/vmware-casa/bin/Netcheck.py
```

On each node, you are presented with a list of attempted connections. If a node cannot connect to the required port, it is reported in the list. Ports that do not connect must be investigated.

Note Only one port is required within the range of 10002-10010 and 20002-20010.

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.

After You Log In

After you log in to vRealize Operations Manager from a web browser, you see the Quick Start page. You can set any dashboard to be the landing page instead of the Quick Start page. Click the **Actions** menu on a dashboard that you want to set as the landing page and select **Set as Home landing page**. To remove the dashboard as the home landing page, click the **Actions** menu on the relevant dashboard and select **Reset from Home landing page**.

The Quick Start page provides an overview of key areas of vRealize Operations Manager.

Quick Start Page Before Cloud Accounts Are Configured

When you log in to vRealize Operations Manager and no cloud accounts are configured, the Quick Start page displays guided tours in the Optimize Performance, Optimize Capacity, Troubleshoot, and Manage Configuration sections. Watch these guided tours to understand how the product functions. If your user account does not have administrative rights, then the Quick Start page prompts you to contact the administrator for configuration of cloud accounts.

If you have logged in using an administrative account, you must set the currency in the **Global Settings** page. In the menu, click **Administration**, and then in the left pane click **Management > Global Settings**. You can do so from the message that you see in the Quick Start page when you log in for the first time. Optionally, you can close the message. Once you set a currency, you cannot change it. As an administrator, you must also first set up a cloud account or configure an adapter before you can start using vRealize Operations Manager. Until you do so, you see links to guided tours about vRealize Operations Manager.

A new license key is required for vRealize Operations Manager 7.0 and later versions. All license keys except vSOM Enterprise Plus and its add-ons are invalidated. The product works in evaluation mode until a new valid license key, which can be obtained from the [MyVMware](#) portal, is installed. After login, if you see the "You are using an evaluation license. Please consider applying a new license by the end of the evaluation period." message in the Quick Start page, you must add a new license before the end of the 60-day evaluation period in the Licensing page. To add a new license, from the message, click **Actions > Go to Licensing**.

Note If you added new licenses when you upgraded to vRealize Operations Manager 7.0, you can skip this step.

After logging in, if you see a message like, "vRealize Operations Manager internal certificates will expire on dd/mm/yyyy. Please install a new certificate before the expiry date. For details, see KB 71018" in the Quick Start page, you must upgrade your internal certificates for vRealize Operations Manager using the certificate renewal PAK file from the vRealize Operations Manager Administrator interface. For more information, see the following KB article [71018](#).

Quick Start Page After Cloud Accounts Are Configured

When you log in to vRealize Operations Manager after the cloud accounts or adapter instances are configured, and the initial setup is complete, the Quick Start displays the following sections.

Optimize Performance

Displays links to workload optimization, right sizing, recommendations, and optimization history.

Optimize Capacity

Displays links to assess capacity, reclaim resources, plan scenarios and assess costs.

Troubleshoot

Displays links to the troubleshooting workbench, alerts, logs, and dashboards.

Manage Configuration

Displays links to the compliance page. Links to the dashboard that displays the configuration of your virtual machines.

Click **View More** to access the following sections:

Extend Monitoring

Displays links to the following VMware website:

- VMware SDDC Health Monitoring Solution
- vRealize Operations Aggregator Management Pack 2.0

Learn and Evaluate

Displays links to the vRealize Operations Guided Tour, Evaluate vRealize Suite, Additional Learning, and Evaluate Sample Dashboards websites.

Run Assessments

Displays shortcut links to the VMware vRealize Cloud Management Assessment and vSphere Optimization Assessment (Deprecated) pages in vRealize Operations Manager.

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 Log in to a vCenter Server system using a vSphere Web Client and select a vCenter Server instance in the vSphere Web Client navigator.**a** Find the **Virtual Machine** in the hierarchy and click **Launch Console**.

Note You can also use the vSphere Client to launch the node console by direct access after enabling the SSHD service.

The virtual machine console opens in a new tab of the Web browser.

2 Locate the node console and click **Launch Console**.**3** In vCenter, use Alt+F1 to access the login prompt and 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.**4** 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 primary node.
- 2 Enter the username **admin** and the password that you defined when you configured the primary 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.

Upgrade, Backup and Restore

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. In case modifications are required, you can manually update the hosts file after completing the software update.

To download the PAK file for vRealize Operations Manager, go to [Download VMware vRealize Operations](#) page.

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. For more information, see [Creating a Backup and Importing Content](#) in the [Managing Content](#) chapter of the [Configuration Guide](#).

Back Up and Restore

Back up and restore your vRealize Operations Manager system regularly 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 back up 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 back up and restore vRealize Suite components by using vSphere Data Protection and NetBackup, see the [Back up and Restore](#) section in the [vRealize Suite Information Center](#).

It is highly recommended to take a backup during quiet periods. Since a snapshot based backup happens at the block level, it is important that there are limited or no changes being performed by a user on the cluster configuration. This will ensure that you have a healthy backup.

It is best to take the cluster offline before you back up the vRealize Operations Manager nodes. This will ensure the data consistency across the nodes and internally in the node. You can either shut down the VM before the backup or enable quiescing.

If the cluster remains online, backup your vRealize Operations Manager multi-node cluster by using vSphere Data Protection or other backup tools, disable quiescing of the file system.

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 3-7. 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. Primary 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**.
- Since 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 pane.
- 3 Click **Install a Software Update** in the main pane.

- 4 Follow the steps in the wizard to locate and install your PAK file.

This updates the OS on the virtual appliance and restarts each virtual machine.

Note When you upgrade to vRealize Operations Manager 8.3 version from a version prior to 8.0, the base OS automatically changes to Photon. Any customization done to the OS, for example, files or directories created somewhere on the root partition, like `~/.ssh/authorized_keys` of the vRealize Operations Manager appliance gets deleted after the upgrade.

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.

Note When you update vRealize Operations Manager to a latest version, all nodes get upgraded by default.

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 user name 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.

Before Upgrading to vRealize Operations Manager 8.3

With every vRealize Operations Manager release, many metrics are either discontinued or disabled. These changes update the capacity analytics and improve the product scale. VMware has made many of these changes transparent or nearly so. Still, multiple changes can impact management packs that you might be using, along with the dashboards and reports that you have created. Therefore, before upgrading, run the vRealize Operations Manager Pre-upgrade Readiness Assessment Tool (Assessment Tool) that helps you understand the precise impact on your environment through a detailed report.

Why Run the Assessment Tool

Various changes in vRealize Operations Manager can impact the user experience. When you run the Assessment Tool, you get an HTML-formatted report identifying all the points in your system affected by the changes. Further, the Assessment Tool gives recommendations for the correct changes to be made in your content for when you upgrade from a previous release.

Note You must run the Assessment Tool on the instance of the vRealize Operations Manager installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the vRealize Operations Manager Administration user interface.

The Assessment tool validates your environment to ensure it is ready for the upgrade. For example, if the ESXi version does not match the product requirements, the assessment tool will identify the issue and provide you with a recommendation in the Systems Validation tab.

For detailed instructions on running the Assessment Tool, see [Running the vRealize Operations Manager 8.3 Pre-Upgrade Readiness Assessment Tool](#).

To view the upgrade path from an earlier version of vRealize Operations Manager to 8.3, see [vRealize Operations Manager Upgrade Path](#).

Running the vRealize Operations Manager 8.3 Pre-Upgrade Readiness Assessment Tool

Before upgrading, you can gauge the impact on your system by running the vRealize Operations Manager Pre-Upgrade Readiness Assessment Tool (Assessment Tool). The tool generates a report detailing the precise impact on your environment and gives suggestions for replacement metrics.

Using the Assessment Tool consists of four distinct steps:

- 1 Download the PAK file from <https://my.vmware.com/group/vmware/get-download?downloadGroup=VROPS-830>.
- 2 Run the vRealize Operations Manager Pre-Upgrade Readiness Assessment Tool.
- 3 Extract the report from the generated ZIP file.
- 4 Click the various items in the report to link to the solutions grid.

Note You must run the Assessment Tool on the instance of the vRealize Operations Manager installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the vRealize Operations Manager Administration user interface.

Prerequisites

You must have administrator privileges in your current installation of vRealize Operations Manager to download and run the Assessment Tool. For more information on using the upgrade assessment tool, see the following KB article [67311](#).

Procedure

- 1 Download the Assessment Tool PAK from <https://my.vmware.com/group/vmware/get-download?downloadGroup=VROPS-830> to your local machine. Search for APUAT or vRealize Operations - Upgrade Assessment Tool.

- 2 Open a browser and navigate to the vRealize Operations Manager administrator console: `https://<master_node_IP>/admin`.

Then log into the administrator user interface with the user ID **admin** and the associated password.

- 3 In the left pane of the administration home page, click **Software Update**.

The Software Update screen appears.

- 4 Click **Install a Software Update** at the top of the screen.

The Add Software Update workspace appears.

- 5 Click the **Browse** link and navigate to the PAK file you downloaded in Step 1.

A check mark appears next to the statement: **The selected file is ready to upload and install. Click UPLOAD to continue.**

- 6 Ensure that a check mark appears next to the statement: **Install the PAK file even if it is already installed.**

Leave blank the check box next to Reset Default Content...

- 7 Click the **UPLOAD** link.

The PAK file is uploaded from your local machine to vRealize Operations Manager. Uploading may take a few minutes.

- 8 Once the PAK file is uploaded, click **NEXT**.

The End User License Agreement appears.

- 9 Click the check box next to the statement: **I accept the terms of this agreement.**

Click **NEXT**. The Important Update and Release Information screen appears.

- 10 Review the release information and click **NEXT**. At the Install Software Update screen, click **INSTALL**.

The Software Update screen appears again, this time with a rotating icon and an **installation in progress...** bar marking the progress of the PAK file and assessment as they run on your environment. The process can take from five to 20 minutes, depending on the size of your system.

- 11 When the process is complete, click **Support** in the left pane.

The Support screen appears.

- 12 Select the **Support Bundles** option above the toolbar.

The available support bundles are listed.

- 13 Locate the support bundle most recently created. Click the chevron next to the bundle name to open the file and select it, then click the download link on the toolbar to save the support bundle ZIP file to your local files.

- 14 To review the report, extract the files from the ZIP file and open the HTML file. (Do not open the CSV file, it is for VMware use only.)

The report is a graphical depiction of your vRealize Operations Manager UI components - dashboards, reports, management packs, alerts, heat maps, and so on - and includes the number of deprecated metrics impacting each component. For example, you might find that 10 of your 25 dashboards contain a total of 15 deprecated metrics.

- 15 Click a component.

The report details for that component are listed following the graphics, under Impacted Component Details. Taking dashboards as an example, the list provides - for each dashboard - the dashboard name, owner, widgets removed, metric-impacted views, and metric-impacted widgets. The deprecated metrics are live links.

- 16 Click a live metric link.

A browser window opens at URL <http://partnerweb.vmware.com/programs/vrops/DeprecatedContent.html> with the selected metric highlighted in a table of like metrics. If a replacement metric is available for the deprecated metric, it is listed in the same row by name and metric key. You might choose to install the new metric in place of the deprecated metric.

- 17 Repeat Steps 15 and 16 for all your components.

If you replace the deprecated metrics with new metrics, or update each component to provide needed information without the deprecated metrics, your system is ready for the upgrade.

- 18 Rerun the entire assessment process from Step 1 to confirm that your system is no longer impacted or at least mostly not impacted by the metrics changes.

- 19 Once you have upgraded to vRealize Operations Manager 8.3, fix the remaining issues with replacement metrics available in the new release.

Results

Your vRealize Operations Manager components are updated to work correctly in the 8.3 release.

What to do next

Once you have installed vRealize Operations Manager 8.3, conduct, at a minimum, random testing to determine if system metrics are operating as you expect. Monitor the platform on an ongoing basis to confirm that you are receiving the correct data.

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 Optimization](#)

Connecting vRealize Operations Manager to Data Sources

You can extend the monitoring capabilities of vRealize Operations Manager by installing and configuring management packs in vRealize Operations Manager to connect to and analyze data from external data sources in your environment. Once connected, you use vRealize Operations Manager to monitor and manage objects in your environment.

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

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

VMware solutions include adapters for Storage Devices, Log Insight, NSX for vSphere, Network Devices, and VCM. Third-party solutions include AWS, SCOM, EMC Smarts, and many others.

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

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

- VMware vSphere
- VMware vRealize Log Insight

- VMware vRealize Assessments
- VMware vSAN
- vRealize Operations Service Discovery Management Pack
- VMware vRealize Automation 8.x
- VMware Management Pack for AWS
- VMware Management Pack for Microsoft Azure
- VMware vRealize Operations Management Pack for NSX-T
- VMware vRealize Operations Management Pack for VMware Cloud on AWS
- VMware vRealize Network Insight

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

- Operating Systems/Remote Service Monitoring
- VMware vRealize Application Management Pack
- VMware vRealize Compliance Pack for PCI
- VMware vRealize Compliance Pack for ISO
- VMware vRealize Compliance Pack for HIPAA
- VMware vRealize Compliance Pack for FISMA
- VMware vRealize Compliance Pack for CIS
- VMware vRealize Compliance Pack for DISA
- VMware vRealize Ping

Upgrade Considerations

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

If you upgrade from an earlier version of vRealize Operations Manager, your management pack files are copied to the `/usr/lib/vmware-vcops/user/plugins/.backup` file in a folder with the date and time as the folder name. Before migrating your data to your new vRealize Operations Manager instance, you must configure the adapter instances again. If you have customized the adapter, your adapter customizations are not included in the migration, and you must reconfigure the customizations.

If you update a management pack in vRealize Operations Manager to a newer version, and you have customized the adapter, your adapter customizations are not included in the upgrade, and you must reconfigure them.

Solutions Repository

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

Where You Find the Repository Page

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

Table 4-1. Repository Page Options

Options	Descriptions
VMware Native Management Packs	
Name	Name of the solution.
Activate	<p>Installs the native management pack. You can configure cloud management packs after activation from Solutions > Cloud Accounts. You can configure all other management packs after activation from Solutions > Other Accounts.</p> <p>The activation starts only if all the cluster's nodes are accessible.</p> <p>Note Pre-Installed management packs are activated by default. You can configure them from the Cloud Accounts or the Other Accounts page as applicable. Click Add Account configure the solutions.</p>
Deactivate	<p>Uninstalls the management pack.</p> <p>Note Pre-installed management packs cannot be deactivated.</p>
Status	<p>Indicates whether the management pack has been configured or not. A green tick symbolizes that the management pack has been successfully installed. If configured, you can view the number of accounts associated to it.</p> <p>To view or edit the accounts, click the account link to navigate to the accounts page associated to the management pack.</p>
Provided By	Name of the vendor or manufacturer who created the solution.
Version	Version and build number identifiers of the solution.
View Content	Displays the list of content that has been deployed using the management pack.

Table 4-1. Repository Page Options (continued)

Options	Descriptions
Reset Default Content	<p>This option is only available for the VMware vSphere solution.</p> <p>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.</p> <p>When you upgrade your current version of vRealize Operations Manager, you must select this option to overwrite alert definitions and symptom definitions. If you do not overwrite alert and symptom definitions, compliance rules use a mixture of new and outdated definitions.</p>
Other Management Packs	
Add/Upgrade	You can add a management pack. For details, see Adding Solutions .

Managing Solutions in vRealize Operations Manager

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

How Solutions Work

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

Where You Find Solutions

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

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

Data Collection Notifications

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

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

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

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

Failed Solution Installation

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

Manage Cloud Accounts

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

The Cloud Accounts page includes a toolbar of options.

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

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

Table 4-2. Cloud Accounts Grid Options

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

Table 4-2. Cloud Accounts Grid Options (continued)

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

Manage the Other Solutions

To add and configure the other solutions, see [Adding Other Accounts](#)

Adding Cloud Accounts

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

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

To manage accounts for the vSphere solution, see [Cloud Account Information - VMware vSphere Account Options](#).

To add and configure accounts for the Management Pack for AWS, see [Add a Cloud Account for Management Pack for AWS](#)

To add and configure accounts for the Management Pack for Microsoft Azure, see [Add a Cloud Account for the Management Pack for Microsoft Azure](#).

To add and configure accounts for VMware Cloud on AWS, see [Configuring a VMware Cloud on AWS Instance in vRealize Operations Manager](#)

Prerequisites

Note

- Activate the cloud account before adding and configuring cloud accounts.
- VMware vSphere solution is activated by default and cannot be deactivated.

Importing Cloud Accounts

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

Prerequisites

- Verify that vRealize Automation 8.x is enabled from **Administration > Management > Integrations** in vRealize Operations Manager .
- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data.
- Verify that the user has privileges of Organizational Owner and Cloud Assembly administrator set in vRealize Automation.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Cloud Accounts > Import Accounts**.
- 2 From the **Import Accounts** page, select the cloud account you want to import.
- 3 To override an existing credential from vRealize Automation, click the **Edit** icon next to **Edit Credential**.
 - Select the existing credential from the **Credential** drop-down menu and click **Save**.
 - To add a new credential, click the plus icon next to the **Credential** drop-down menu and enter the credential details and click **Save**.
- 4 Select the collector/group from the drop-down menu.
- 5 Click **Validate** to verify that the connection is successful.
- 6 Click **Import**.

Results

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

Manage Other Accounts

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

Note You must activate solutions before configuring them. For more information, see [Solutions Repository](#)

The Other Accounts page includes a toolbar of options.

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

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

Table 4-3. Cloud Accounts Grid Options

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

Manage the Cloud Solutions

To add and configure the cloud accounts, see [Manage Other Accounts](#)

Adding Other Accounts

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

Note

- Activate the solutions before adding and configuring other accounts.

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

The options available depend on the selected solution.

Configuring VMware vRealize Ping Adapter Instances

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

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

Procedure

- 1 In the menu, click **Administration**, and then from the left pane click **Solutions > Other Accounts > Add Accounts**.
- 2 Click the VMware vRealize Ping adapter instance.
- 3 Configure the VMware vRealize Ping adapter instance.

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

- 4 Click **Save**.

Results

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

Adding Solutions

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

How Added Solutions Work

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

Where You Add Solutions

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

Add Solutions Wizard Options

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

Before you install the PAK file, or upgrade your vRealize Operations Manager instance, clone any customized content to preserve it. Customized content can include alert definitions, symptom definitions, recommendations, and views.

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

Table 4-4. Wizard Options

Option	Description
Page 1	
Browse a Solution	Navigate to your copy of a management pack PAK file.
Upload	To prepare for installation, copy the PAK file to vRealize Operations Manager.
Install the PAK file even if it is already installed	If the PAK file was already uploaded, reload the PAK file using the current file, but leave user customizations in place. Do not overwrite or update the solution alerts, symptoms, recommendations, and policies.
Reset Default Content	<p>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.</p> <p>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.</p>
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	<p>Read and agree to the end-user license agreement.</p> <p>Note Click Next to install the solution. The installation starts only if all the cluster's nodes are accessible.</p>
Page 3	
Installation Details	Review the installation progress, including the vRealize Operations Manager nodes where the adapter was installed.

Manage Integrations

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

Where You Find Integrations

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

Table 4-5. Integration Page Options

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

Managing Solution Credentials

Credentials are the user accounts that vRealize Operations Manager uses to enable one or more solutions and associated adapters, and to establish communication with the target data sources. The credentials are supplied when you configure each adapter. You can add or modify the credential settings outside the adapter configuration process to accommodate changes to your environment.

For example, if you are modifying credentials to accommodate changes based on your password policy, the adapters configured with these credentials begin using the new user name and password to communicate between vRealize Operations Manager and the target system.

Another use of credential management is to remove misconfigured credentials. If you delete valid credentials that were in active use by an adapter, you disable the communication between the two systems.

If you need to change the configured credential to accommodate changes in your environment, you can edit the credential settings without being required to configure a new adapter instance for the target system. To edit credential settings, click **Administration** on the menu, and in the left pane, click **Management> Credentials**.

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

Credentials

The credentials are the collection configuration settings, for example, user names and passwords, that the adapters use to authenticate the connection on the external data sources. Other credentials can include values such as domain names, pass phrases, or proxy credentials. You can configure for one or more solutions to connect to data sources as you manage your changing environment.

Where You Find Credentials

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

Table 4-6. Credentials Options

Option	Description
Toolbar options	<p>Manages the selected credential.</p> <ul style="list-style-type: none"> ■ Add. Add new credentials for an adapter type that you can later apply when configuring an adapter. ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Modify the selected credentials, usually when the user name and password require a change. The change is applied to the current adapter credentials and the data source continues to communicate with vRealize Operations Manager. ■ Delete . Remove the selected credentials from vRealize Operations Manager. If you have an adapter that uses these credentials, the communication fails and you cease monitoring the objects that the adapter was configured to manage. Commonly used to delete misconfigured credentials.
Filtering options	Limits the displayed credentials based on the adapter or credential types.
Credential name	Description of user-defined name that you provide to manage the credentials. Not the account user name.
Adapter Type	Adapter type for which the credentials are configured.
Credential Type	Type of credentials associated with the adapter. Some adapters support multiple types of credentials. For example, one type might define a user name and password, and another might define a pass code and key phrase.

Manage Credentials

To configure or reconfigure credentials that you use to enable an adapter instance, you must provide the collection configuration settings, for example, user name and password, that are valid on the target system. You can also modify the connection settings for an existing credential instance.

Where You Manage Credentials

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

Manage Credentials Options

The Manage Credentials dialog box is used to add new or modifies existing adapter credentials. The dialog box varies depending on the type of adapter and whether you are adding or editing. The following options describe the basic options. Depending on the solution, the options other than the basic ones vary.

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

Table 4-7. Manage Credential Add or Edit Options

Option	Description
Adapter Type	Adapter type for which you are configuring the credentials.
Credential Kind	Credentials associated with the adapter. The combination of adapter and credential type affects the additional configuration options.
Credential Name	Descriptive name by which you are managing the credentials.
User Name	User account credentials that are used in the adapter configuration to connect vRealize Operations Manager to the target system.
Password	Password for the provided credentials.

Managing Collector Groups

vRealize Operations Manager uses collectors to manage adapter processes such as gathering metrics from objects. You can select a collector or a collector group when configuring an adapter instance.

If there are remote collectors in your environment, you can create a collector group, and add remote collectors to the group. When you assign an adapter to a collector group, the adapter can use any collector in the group. Use collector groups to achieve adapter resiliency in cases where the collector experiences network interruption or becomes unavailable. If this occurs, and the collector is part of a group, the total workload is redistributed among all the collectors in the group, reducing the workload on each collector.

Collector Group Workspace

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

Rebalancing an Adapter Instance

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

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

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

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

Where You Manage Collector Groups

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

Table 4-8. Collector Group Summary Grid

Options	Description
Collector Group toolbar	<p>To manage collector groups, use the toolbar icons.</p> <ul style="list-style-type: none"> ■ Add. Add a collector group ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Modify the collector group by adding or removing remote collectors. ■ Delete. Remove the selected collector group. ■ Rebalance collector group. Rebalance one collector group at a time. If you have permissions to manage clusters, you can rebalance the workload across the collectors and the remote collectors in the collector group. The rebalance action moves objects from one collector group to another to rebalance the number of objects on each collector in the collector group. If a disk rebalance is already in progress, the collector rebalance does not run.
Collector Group Name	The name given to the collector group when the collector group is created.
Description	Description given to the collector group when the collector group is created.
All Filters	Displays the list of collector groups in the summary grid by collector group name, description, collector name, or IP address.
Quick Filter Name	Filters the list of collector groups according to the name of the collector group entered.

Table 4-9. Collector Group Details Grid

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

Adding a Collector Group

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

Where You Add New Collector Groups

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

Add New Collector Group Workspace

Option	Description
Name	Name of the collector group.
Description	Description of the collector group.

Option	Description
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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ Collector Name ■ IP Address ■ Status

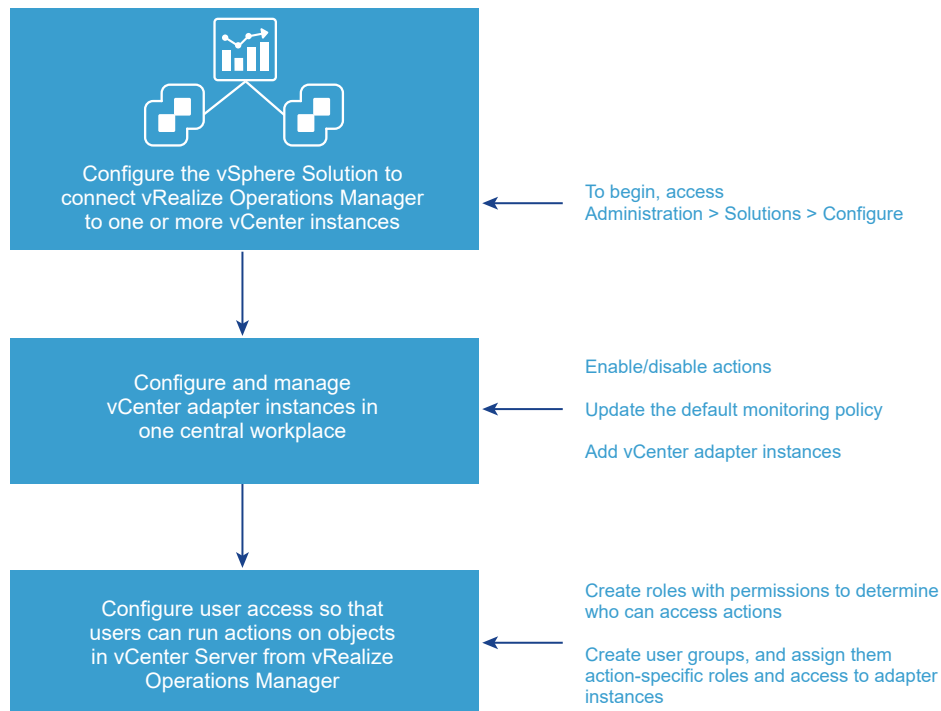
VMware vSphere Solution in vRealize Operations Manager

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

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

Configuring the vSphere Solution

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



How Adapter Credentials Work

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

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

Controlling User Access to Actions

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

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

To configure a vCenter Server cloud account, see [Configure a vCenter Server Cloud Account in vRealize Operations Manager](#).

Configure a vCenter Server Cloud Account in vRealize Operations Manager

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

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

Prerequisites

- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data, see [Privileges Required for Configuring a vCenter Adapter Instance](#). If the provided credentials have limited access to objects in vCenter Server, all users, regardless of their vCenter Server privileges see only the objects that the provided credentials can access. At a minimum, the user account must have Read privileges and the Read privileges must be assigned at the data center or vCenter Server level.

Procedure

- 1 On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**.
- 2 On the Cloud Accounts page, click **Add Accounts**.
- 3 On the Accounts Type page, click **vCenter**.
- 4 Enter a display name and description for the cloud account.
 - Display name. Enter the name for the vCenter Server instance as you want it to appear in vRealize Operations Manager. A common practice is to include the IP address so that you can readily identify and differentiate between instances.
 - Description. Enter any additional information that helps you manage your instances.

- 5 In the vCenter Server text box, enter the FQDN or IP address of the vCenter Server instance to which you are connecting.

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

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

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

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

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

- 7 Determine which vRealize Operations Manager collector or collector group is used to manage the cloud account. If you have only one cloud account, select **Default collector group**. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.
- 8 The cloud account is configured to run actions on objects in the vCenter Server from vRealize Operations Manager. If you do not want to run actions, deselect **Enable** for Operational Actions.
- 9 Click **Validate Connection** to validate the connection with your vCenter Server instance.
- 10 In the **Review and Accept Certificate** dialog box, review the certificate information.
 - ◆ If the certificate presented in the dialog box matches the certificate for your target vCenter Server, click **OK**.
 - ◆ If you do not recognize the certificate as valid, click **Cancel**. The test fails and the connection to vCenter Server is not completed. You must provide a valid vCenter Server URL or verify the certificate on the vCenter Server is valid before completing the adapter configuration.
- 11 To modify the advanced options regarding collectors, object discovery, or change events, expand the **Advanced Settings**.

For information about these advanced settings, see [Cloud Account Information - VMware vSphere Account Options](#).

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

For information about monitoring goals, see [Cloud Account Information - VMware vSphere Account Options](#).

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

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

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

- ◆ If the vCenter Server already has a vRealize Operations Manager instance registered to it, you can override the existing registrations with your instance of vRealize Operations Manager. Click **Yes** to replace the existing registration with your vRealize Operations Manager instance.
- ◆ To proceed with the configuration without registering your vRealize Operations Manager, click **No**.

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

Results

The cloud account is added to the list. vRealize Operations Manager begins collecting metrics, properties, and events from the vCenter Server instance. Depending on the number of managed objects, the initial collection can take more than one collection cycle. A standard collection cycle begins every five minutes.

For information about the network port that vRealize Operations Manager uses to communicate with a vCenter Server system and vRealize Operations Manager components, see <http://ports.vmware.com>.

What to do next

You can enable vSAN Configuration for your cloud account. For more information, see [Configure a vSAN Adapter Instance](#).

You can use the vCenter Server for service discovery, see [Configure Service Discovery](#).

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

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

The Register vCenter Server dialog box appears.

- 2 Click the **Use collection credentials** check box.
 - Click **Unregister** to remove any existing registrations.

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

Privileges Required for Configuring a vCenter Adapter Instance

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

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

Table 4-10. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection

Task	Privilege
Property Collection	System > Anonymous Note When you add a custom role and do not assign any privileges to it, the role is created as a Read Only role with three system-defined privileges: System.Anonymous , System.View , and System.Read . See, Using Roles to Assign Privileges .
Objects Discovery Events Collection	Profile-Driven Storage > View Storage views > View Profile-Driven Storage > Profile-Driven Storage View Datastore > Browse Datastore System > View Note This permission is provided with the Read-Only role.
Performance Metrics Collection	Performance > Modify intervals System > Read Note This permission is provided with the Read-Only role.
Service Discovery	Virtual Machine > Guest Operations > Guest Operation alias modification Virtual Machine > Guest Operations > Guest Operation alias query Virtual Machine > Guest Operations > Guest Operation modifications Virtual Machine > Guest Operations > Guest Operation program execution Virtual Machine > Guest Operations > Guest Operation queries

Table 4-10. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection (continued)

Task	Privilege
Tag Collection	Global > Global tag Global > Global health Global > Manage custom attributes Note This privilege is required only if the tags are associated with custom attributes. Global > System tag Global > Set custom attribute
Monitor the Namespace Resource Pool or objects in the Resource Pool.	The account for the adapter instance also needs to be a member of <code>Administrators@vsphere.local</code> on the vCenter Server.

Table 4-11. Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions

Task	Privilege
Set CPU Count for VM	Virtual Machine > Configuration > Change CPU Count
Set CPU Resources for VM	Virtual Machine > Configuration > Change Resource
Set Memory for VM	Virtual Machine > Configuration > Change Memory
Set Memory Resources for VM	Virtual Machine > Configuration > Change Resource
Delete Idle VM	Virtual machine > Edit Inventory > Remove
Delete Powered Off VM	Virtual machine > Edit Inventory > Remove
Create Snapshot for VM	Virtual Machine > Snapshot Management > Create Snapshot
Delete Unused Snapshots for Datastore	Virtual Machine > Snapshot Management > Remove Snapshot
Delete Unused Snapshot for VM	Virtual Machine > Snapshot Management > Remove Snapshot
Power Off VM	Virtual Machine > Interaction > Power Off
Power On VM	Virtual Machine > Interaction > Power On
Shut Down Guest OS for VM	Virtual Machine > Interaction > Power Off
Move VM	<ul style="list-style-type: none"> ■ Resource > Assign Virtual Machine to Resource Pool ■ Resource > Migrate Powered Off Virtual Machine ■ Resource > Migrate Powered On Virtual Machine ■ Datastore > Allocate Space Note Combining these four permissions allows the service account to perform Storage vMotion and regular vMotion of an object therefore allowing vRealize Operations Manager to perform the given operations.

Table 4-11. Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions (continued)

Task	Privilege
Optimize Container	<ul style="list-style-type: none"> ■ Resource > Assign Virtual Machine to Resource Pool ■ Resource > Migrate Powered Off Virtual Machine ■ Resource > Migrate Powered On Virtual Machine ■ Datastore > Allocate Space
Schedule Optimize Container	<ul style="list-style-type: none"> ■ Resource > Assign Virtual Machine to Resource Pool ■ Resource > Migrate Powered Off Virtual Machine ■ Resource > Migrate Powered On Virtual Machine ■ Datastore > Allocate Space
Set DRS Automation	Host > Inventory > Modify Cluster
Provide data to vSphere Predictive DRS	External stats provider > Update External stats provider > Register External stats provider > Unregister

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

Configure User Access for Actions

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

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

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

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Access > Access Control**.
- 2 To create a role:
 - a Click the **Roles** tab.
 - b Click the **Add** icon, and enter a name and description for the role.
- 3 To apply permissions to the role, select the role, and in the Permissions pane, click the **Edit** icon.
 - a Expand **Environment**, and then expand **Action**.
 - b Select one or more of the actions, and click **Update**.

4 To create a user group:

- a Click the **User Groups** tab, and click the **Add** icon.
- b Enter a name for the group and a description, and click **Next**.
- c Assign users to the group, and click the **Objects** tab.
- d Select a role that has been created with permissions to run actions, and select the **Assign this role to the user** check box.
- e Configure the object privileges by selecting each adapter instance to which the group needs access to run actions.
- f Click **Finish**.

What to do next

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

Cloud Account Information - VMware vSphere Account Options

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

Where You Find the Solution - VMware vSphere

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

Account Information - VMware vSphere Account Options

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

Table 4-12. Advanced Settings Options

Option	Description
Advanced Settings	Provides options related to designating specific collectors to manage this cloud account, managing object discovery and change events.
Auto Discovery	<p>Determines whether new objects added to the monitored system are discovered and added to vRealize Operations Manager after the initial configuration of the cloud account.</p> <ul style="list-style-type: none"> ■ If the value is true, vRealize Operations Manager collects the information about any new objects that are added to the monitored system after the initial configuration. For example, if you add more hosts and virtual machines, these objects are added during the next collections cycle. This is the default value. ■ If the value is false, vRealize Operations Manager monitors only those objects that are present on the target system when you configure the cloud account.

Table 4-12. Advanced Settings Options (continued)

Option	Description
Process Change Events	<p>Determines whether the cloud account uses an event collector to collect and process the events generated in the vCenter Server instance.</p> <ul style="list-style-type: none"> ■ 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	When set to false, reduces the collected data set by omitting collection of the associated category.
Enable Collecting Virtual Machine Folder	
Enable Collecting vSphere Distributed Port Group	
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	<p>Reduces the collected data set by limiting the number of virtual machine collections.</p> <p>To omit data on virtual machines and have vRealize Operations Manager collect only host data, set the value to zero.</p>
Provide data to vSphere Predictive DRS	<p>vSphere Predictive DRS proactively load balances a vCenter Server cluster to accommodate predictable patterns in the cluster workload.</p> <p>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.</p> <p>Predictive DRS must also be enabled for the Compute Clusters managed by the vCenter Server instances monitored by vRealize Operations Manager. Refer to the <i>vSphere Resource Management Guide</i> for details on enabling Predictive DRS on a per Compute Cluster basis.</p> <p>When set to true, designates vRealize Operations Manager as a predictive data provider, and sends predictive data to the vCenter Server. You can only register a single active Predictive DRS data provider with a vCenter Server at a time.</p>
Enable Actions	Enabling this option helps in triggering the actions that are related to vCenter.
Cloud Type	Provides an ability to identify the type of vCenter is used in vRealize Operations Manager. By default, the cloud type is set to Private Cloud.
vCenter ID	A globally unique identifier associated with the vCenter Server instance.
Disable collecting Guest File Systems with names containing	Provide comma separated list of strings. If these strings are found in any guest files system mount point name, that guest file system will not be collected.
Collection Interval (Minutes)	The interval between collection of data from the vCenter Server.
Dynamic Thresholding	This setting is enabled by default.

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

Table 4-13. Define Monitoring Goals Page Options

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

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

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

VMware Cloud on AWS

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

Configuring a VMware Cloud on AWS Instance in vRealize Operations Manager

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

Prerequisites

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

Procedure

- 1 On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**.
- 2 On the Cloud Accounts page, click **Add Accounts**.
- 3 On the Accounts Type page, click **VMware Cloud on AWS**.
- 4 Enter a display name and description for the cloud account.
 - Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in vRealize Operations Manager.

- Description. Enter any additional information that helps you manage your instances.
- 5 To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.
 - Credential Name. The name by which you are identifying the configured credentials.
 - CSP Refresh Token. A CSP API token. For details on generating an API token, see [Generating CSP API Token](#).
 - Proxy Host
 - Proxy Port
 - Proxy username
 - Proxy Password
 - Proxy Domain
 - 6 Determine which vRealize Operations Manager collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

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

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

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

- 8 Click **Validate Connection** to validate the connection.
- 9 You can monitor the costs of running your VMware Cloud on AWS infrastructure by bringing in the billing from VMware Cloud on AWS to vRealize Operations Manager. To do so, enable the costing option in **Advanced Settings**.
- 10 Click **Save**.

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

- 11 Click **Configure**.

- 12 Configure the vCenter adapter:

- a Click the **Add** icon, and enter the required credentials.
 - Credential Name. The name by which you are identifying the configured credentials.
 - User Name. The vCenter user name.
 - Password. The vCenter password configured for that vCenter user name.
- b Select the required collector group.
- c Click **Next**.

13 Configure the vSAN adapter.

- a Enable the **vSAN configuration** option.
- b Click **Validate Connection** to validate the connection.
- c Click **Next**.

14 Configure the NSX-T adapter.

- a Enable the **NSX-T configuration** option.
- b Click **Validate Connection** to validate the connection.
- c Click **Next**.

15 Click **Save This SDDC**.

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

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

Known Limitations

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

- Setting the Costing Enabled option to true, collects VMware Cloud on AWS bills. The cost collection depends on the invoice generation. For instances, where the invoices are not generated, contact the VMware on Cloud for AWS support.
- vRealize Operations Manager does not calculate the datacenter, cluster, or VM costs for VMware on AWS SDDCs even if the bills are being collected. Cost drivers also cannot be configured for VMware Cloud on AWS.
- The `cloudadmin@vmc.local` user in VMware Cloud has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines on VMware Cloud. Active and consumed memory utilizations continue to work in this case.
- The compliance workflows in vRealize Operations work for virtual machines running on a vCenter Server in VMware Cloud on AWS. The compliance checks for VMware management objects such as Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent does not work because VMware manages cluster configurations.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with vRealize Operations Manager 8.0 onwards. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud does not support vRealize Operations Manager plugin.

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

Generating CSP API Token

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

Prerequisites

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

Procedure

- 1 Log in to the [VMware Cloud Services](#), select your user profile in the top-right corner, and click **My Account**.
- 2 In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
- 3 Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
- 4 Click **Generate**.
- 5 Copy or save the generated token.

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

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

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

Table 4-14. Object Types that NSX-T Discovers

Object Type	Description
NSX-T Adapter Instance	The vRealize Operations management pack for the NSX-T instance.
Logical Switch	Logical segments in the NSX-T environment.
Logical Switches	Group of the logical segments.
Firewall Section	Firewall sections in the NSX-T environment.

Table 4-14. Object Types that NSX-T Discovers (continued)

Object Type	Description
Firewall Sections	Group of firewall sections.
Logical Router	Logical routers in the NSX-T environment.
Logical Routers	Group of tier-0 and tier-1 logical routers.
Tier-0 Routers	Group of tier-0 logical routers.
Tier-1 Routers	Group of tier-1 logical routers.
Group	Groups in the NSX-T environment.
Management Groups	Group of management groups in the NSX-T environment.
Compute Groups	Group of compute groups in the NSX-T environment.
Groups	Group of both management and compute groups.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Inventory**.
- 2 In the list of tags, expand **Adapter Instances** and expand **NSX-T Adapter Instance**.
- 3 Select the adapter instance name to display the list of objects discovered by your adapter instance.
- 4 Slide the display bar to the right to view the object status.

Object Status	Description
Collection State	If green, the object is connected.
Collection Status	If green, the adapter is retrieving data from the object.

- 5 Deselect the adapter instance name and expand the **Object Types** tag.

Each Object Type name appears with the number of objects of that type in your environment.

Azure VMware Solution

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

Configuring an Azure VMware Solution Instance in vRealize Operations Manager

To monitor Azure VMware Solution instances in vRealize Operations Manager, you must configure a vCenter Server cloud account, a vSAN cloud account, service discovery (optional), and the NSX-T adapter.

Procedure

- 1 Configure a vCenter Server Cloud account. For more information, see [Configure a vCenter Server Cloud Account in vRealize Operations Manager](#).
- 2 Configure a vSAN Adapter instance. For more information, see [Configure a vSAN Adapter Instance](#).
- 3 (Optional) Configure Service Discovery. For more information, see [Configure Service Discovery](#).
- 4 Configure the NSX-T adapter. For more information, see [Configuring the NSX-T Adapter](#).

After the adapters and cloud accounts are configured, vRealize Operations Manager discovers and monitors the environment that runs on Azure VMware Solution.

Known Limitations

Review the following list of feature limitations of Azure VMware Solution integration.

- Microsoft manages the compliance of Azure VMware Solution hosts. Ignore the compliance alerts for Azure VMware Solution hosts.
- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters might appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation is not supported on Azure VMware Solution. Ignore all the cost metrics.
- The end-user on the vCenter Server on Azure VMware Solution has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations Manager using the credentials of the vCenter Server on Azure VMware Solution.
- The vCenter Server on Azure VMware Solution does not support the vRealize Operations Manager plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.
- Management VMs are hidden from end user visibility while the respective VMDKs are not. As a result, vRealize Operations Manager considers management VMDKs as orphaned and should be ignored.

Google Cloud VMware Engine

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

Configuring a Google Cloud VMware Engine Instance in vRealize Operations Manager

To monitor Google Cloud VMware Engine instances in vRealize Operations Manager, you must configure a vCenter Server cloud account, a vSAN cloud account, service discovery (optional), and the NSX-T adapter.

Procedure

- 1 Configure a vCenter Server Cloud account. For more information, see [Configure a vCenter Server Cloud Account in vRealize Operations Manager](#).
- 2 Configure a vSAN Adapter instance. For more information, see [Configure a vSAN Adapter Instance](#).
- 3 (Optional) Configure Service Discovery. For more information, see [Configure Service Discovery](#).
- 4 Configure the NSX-T adapter. For more information, see [Configuring the NSX-T Adapter](#).

After the adapters and cloud accounts are configured, vRealize Operations Manager discovers and monitors the environment that runs on Google Cloud VMware Engine.

Known Limitations

Review the following list of feature limitations of Google Cloud VMware Engine integration.

- Google manages the compliance of Google Cloud VMware Engine hosts. Ignore the compliance alerts for Google Cloud VMware Engine hosts.
- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters may appear lower than expected and capacity remaining may appear higher than expected.
- Cost calculation is not supported on Google Cloud VMware Engine. Ignore all the cost metrics.
- The end-user on the vCenter Server on Google Cloud VMware Engine has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to vRealize Operations Manager using the credentials of the vCenter Server on Google Cloud VMware Engine.
- The vCenter Server on Google Cloud VMware Engine does not support the vRealize Operations Manager plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.
- Management VMs are hidden from end user visibility while the respective VMDKs are not. As a result, vRealize Operations Manager considers management VMDKs as orphaned and should be ignored.

AWS

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

Introduction to the Management Pack for AWS

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

Supported AWS Services

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

Service	Object	Description
Elastic MapReduce	EMR Job Flow	Enables developers, researchers, analysts, and data scientists to easily process vast amounts of data.
Elastic Load Balancing	Classic Load Balancer	Provides basic load balancing across multiple Amazon EC2 instances and operates at both the request level and connection level. Classic load balancer is intended for applications that are built within the EC2-Classic network.
	Application Load Balancer	Best suited for load balancing of HTTP and HTTPS traffic, this balancer provides advanced request routing targeted at the delivery of modern application architectures, including microservices and containers.
	Network Load Balancer	Best suited for load balancing of TCP traffic where extreme performance is required.
Amazon EC2	Elastic Compute Cloud	Provides resizable computing capacity in the Amazon Web Services cloud.
	Elastic IP	Elastic IP address is a static IPv4 address designed for dynamic cloud computing, which is reachable from the Internet.
	Elastic Network Interface	Provides a logical networking component in a VPC that represents a virtual network card.
	Placement group	When you run a new EC2 instance, the EC2 service attempts to place the instance in such a way that all your instances are spread out across underlying hardware to minimize correlated failures. You can use placement groups to influence the placement of a group of interdependent instances to meet the needs of your workload.
Amazon EC2 Auto Scaling Group		Web service designed to start or stop Elastic Compute Cloud instances, based on user-defined policies, schedules, and health checks.
Amazon Elastic Block Store	EBS volume	Provides block-level storage volumes for use with Amazon Elastic Compute Cloud instances.

Service	Object	Description
Amazon Relational Database Service	RDS DB Instance	Provides familiar SQL databases while automatically managing administrative tasks.
Amazon ElastiCache	ElastiCache Cluster	Amazon ElastiCache allows you to seamlessly set up, run, and scale popular open-Source compatible in-memory datastores in the cloud. Build data-intensive applications or boost the performance of your existing databases by retrieving data from high throughput and low latency in-memory data stores. Amazon ElastiCache is a popular choice for real-time use cases like Caching, Session Stores, Gaming, Geospatial Services, Real-Time Analytics, and Queuing.
	ElastiCache Node	A node is the smallest building block of an Amazon ElastiCache deployment. It is a fixed-size chunk of secure, network-attached RAM. Each node runs the engine that was selected when the cluster or replication group was created or last modified. Each node has its own Domain Name Service (DNS) name and port. Multiple types of ElastiCache nodes are supported, each with varying amounts of associated memory and computational power.
Amazon Simple Queue	SQS Queue	Provides a reliable, highly scalable, hosted queue for storing messages.
Amazon Elastic Container Registry	ECR Container Repository	Fully managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images.
Amazon Elastic Container Service	ECS Cluster	Highly scalable, high-performance container orchestration service that supports Docker containers and allows you to easily run and scale containerized applications on AWS.
Amazon Elastic Kubernetes Service	EKS Cluster	Allows you to use Kubernetes on AWS without needing to install and operate your own Kubernetes control plane.
AWS Lambda	Lambda Function	AWS Lambda lets you run code without provisioning or managing servers.
Amazon DynamoDB	DynamoDB	Fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale.
Amazon DynamoDB Accelerator (DAX)	DynamoDB Accelerator Cluster	Fully managed, highly available, in-memory cache for DynamoDB.
Amazon Redshift	Redshift Cluster	A fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools.
Amazon Virtual Private Cloud	VPC	Lets you provision a logically isolated section of the AWS Cloud where you can run AWS resources in a virtual network that you define.

Service	Object	Description
	Subnet	Provides a range of IP addresses in your VPC. Use it to run the AWS resources into a specified subnet., for example, use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that will not be connected to the Internet.
	Transit Gateway	
	Security Group	A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. When you run an instance in a VPC, you can assign up to five security groups to the instance. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC can be assigned to a different set of security groups.
	NAT Gateway	Use a network address translation (NAT) gateway to enable instances in a private subnet to connect to the Internet or other AWS services, but prevent the Internet from initiating a connection with those instances.
	VPC VPN Connection	Connect your Amazon VPC to remote networks by using a VPN connection.
Amazon CloudFront	CloudFront Distribution	AmazonCloudFront is a global content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to your viewers with low latency and high transfer speeds.
AWS Cloudformation	Cloudformation Stack	AWS CloudFormation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment.
Amazon S3	S3 Bucket	Object storage built to store and retrieve any amount of data from anywhere.
Amazon WorkSpaces	WorkSpaces	Amazon WorkSpaces is a fully managed, secure Desktop-as-a-Service (Daas) solution that runs on AWS.
Amazon Route 53	Route53 Hosted Zone	A hosted zone is a collection of records for a specified domain.
	Route53 Health Checks	To discover the availability of your EC2 instances, a load balancer periodically sends pings, attempts connections, or sends requests to test the EC2 instances.
AWS Elastic Beanstalk	Elastic Beanstalk	Provides the fastest and simplest way to get web applications up and running on AWS. Simply upload your application code and the service automatically handles all the details such as resource provisioning, load balancing, auto-scaling, and monitoring. Elastic Beanstalk is ideal if you have a PHP, Java, Python, Ruby, Node.js, .NET, Go, or Docker web application.
Amazon Elastic File System	EFS	Provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources.

Note All services are created with the following Service Descriptors:

- Account ID
 - Region
 - Service Type
-

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

Charges for AWS Metrics

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

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

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

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

- Turn off auto discovery and use manual discovery. Select only those objects that are critical to your system.
- Subscribe only to specific critical regions or services.
- Use allowlist and denylist filtering to select object import by name.
- Go to the default attribute package for each object. Turn off collection of metrics that are not critical for your system.

View Management Pack for AWS Objects

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

Procedure

- 1 In the left pane of vRealize Operations Manager, click the **Environment** icon.
- 2 In the Environment Overview, under the Inventory Trees, click **AWS Resources by Regions**.
- 3 To view the child objects, expand the regions and then expand the regions per account.

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

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

Configuring the Management Pack for AWS

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

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

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

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

```
elasticmapreduce:DescribeJobFlows
```

Generate Required Access Keys

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

Prerequisites

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

Procedure

- 1 Log in to Amazon Web Services.
- 2 To generate access keys, see the online documentation on the <https://docs.aws.amazon.com/> site.

Complete the following tasks:

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

Configuring IAM Permissions

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

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

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

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

Table 4-15. IAM Permissions

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

Table 4-15. IAM Permissions (continued)

Service	Required	Permissions
EMR	Required if subscribing to the EMR service.	<p>describe*</p> <pre> { "Effect": "Allow", "Action": ["elasticmapreduce:Describe*", "elasticmapreduce:List*", "elasticmapreduce:ViewEventsFromAllClustersInConsole", "s3:GetObject", "s3:ListAllMyBuckets", "s3:ListBucket", "sdb:Select", "cloudwatch:GetMetricStatistics"], "Resource": "*" }</pre>
RDS	Required if subscribing to RDS service.	For the list of permissions, see RDS Read Only Access json .
ElasticCache	Required if subscribing to ElasticCache service.	For the list of permissions, see Elastic Cache Read Only Access json .
SQS	Required if subscribing to SQS service.	For the list of permissions, see SQS Read Only Access json .
Elastic Container Registry		For the list of permissions, see Elastic Container Read Only Access json .
Elastic Container Service		list*
Lambda		For the list of permissions, see Lambda Read Only Access json and refer to the AWS Lambda policy.
DynamoDB		For the list of permissions, see Dynamo DB Read Only Access json .
DAX		<p>describe*</p> <p>list*</p>
Redshift		For the list of permissions, see Redshift Read Only Access json .
Virtual Private Cloud		For the list of permissions, see VPC Read Only Access json .
Cloud Front Distribution		For the list of permissions, see Cloud Front Distribution Read Only Access json .

Table 4-15. IAM Permissions (continued)

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

Update Configuration Settings in the Properties File

The `amazonaws.properties` file provides configuration options.

Table 4-16. Amazon Web Services Property Settings

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

Table 4-16. Amazon Web Services Property Settings (continued)

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

Tagging Groups

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

Table 4-17. Tagging Groups

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

Add a Cloud Account for Management Pack for AWS

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

Prerequisites

- Obtain the Access Key and Secret Key values. See [Generate Required Access Keys](#). These values are not the same as your log in credentials for the Amazon Web Services site.
- Determine the services for which you collect metrics. See, [Supported AWS Services](#)

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

Table 4-18. Amazon Web Services Regions

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

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

Procedure

- 1 On the menu, click **Administration** and in the left pane, click **Solutions > Cloud Accounts**.
- 2 On the Cloud Accounts page, click **Add Accounts**.
- 3 On the Account Types page, click **AWS**.
- 4 Configure the instance settings.

Option	Action
Name	Enter a name for the adapter instance.
Description	Enter a description.
Credential	<p>Add the credentials used to access the AWS environment by clicking the plus sign.</p> <ul style="list-style-type: none"> ■ Enter an instance name for the credential values you are creating. This is not the name of the adapter instance, but a friendly name for the Access Key and Secret Key credential. ■ Enter your Access Key and Secret Key values. ■ Enter any required local proxy information for your network.
Collector / Group	Select the collector upon which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default has been selected for optimal data collecting.

- 5 Click **Test Connection** to validate the connection.
- 6 Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

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

Option	Action
--------	--------

Collect Custom Metrics

Set this option to true if you want to import all the custom metrics from you AWS account.

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

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

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

7 Click **Save Settings**.

What to do next

Make sure that vRealize Operations Manager is collecting data.

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

Microsoft Azure

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

Supported Azure Services

The Management Pack for Microsoft Azure supports the following services.

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

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

Configuring the Management Pack for Microsoft Azure

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

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

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

Note

- You can install and use the management pack only with an enterprise license of vRealize Operations Manager .
- The management pack has a default time granularity based on the services that it monitors. You cannot configure this granularity against the metrics. You can increase the collection interval but you must not decrease it. The default interval is 10 minutes.

Generate a Client Secret

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

Prerequisites

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

Procedure

- 1 Log in to the Microsoft Azure portal.

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

Complete the following tasks:

- a Create an Azure Active Directory application.

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

- b Under **Access Control (IAM) > Add Role Assignment**, select the role you want to assign to the application. The minimum requirement is 'Reader' or above.
- c Generate a client secret for the application.
- d Copy the subscription ID, directory (tenant) ID, application (client) ID, and client secret to use in your cloud account.

Add a Cloud Account for the Management Pack for Microsoft Azure

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

Prerequisites

- If the Management Pack for Microsoft Azure is deactivated, activate it in vRealize Operations Manager . For more information, see [Solutions Repository](#).
- Generate a client secret in the Microsoft Azure portal to use in this configuration. For more information, see [Generate a Client Secret](#).

Procedure

- 1 On the menu, click **Administration**.
- 2 In the left pane, click **Solutions > Cloud Accounts**.
- 3 Click **Add Account** and select **Microsoft Azure**.
- 4 Enter the cloud account information.

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

- 5 Configure the connection.

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

Option	Action
Credential	<p>Add the credentials used to access Microsoft Azure by clicking the plus sign.</p> <ul style="list-style-type: none"> ■ Enter an instance name for the credential values you are creating. This value is not the name of the adapter instance, but a friendly name for the secret credential. ■ Enter your application ID in your Azure Active Directory. ■ Enter the client secret that you generated for your application in the Microsoft Azure portal. ■ Enter any required local proxy information for your network.
Collector/Group	<p>Select the collector upon which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default is selected for optimal data collecting.</p>

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

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

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

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

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

- 8 Click **Add**.

What to do next

Ensure that the vRealize Operations Manager is collecting data.

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

View Objects for the Management Pack for Microsoft Azure

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

Prerequisites

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

Note When you monitor large-scaled Azure end-points (>1000 objects), change the default collection cycle to 15 minutes so that there is enough time to collect data for all the objects from a scaled end-point.

Procedure

- 1 On the menu, click **Environment**.
- 2 In the left pane, under **Environment Overview**, expand **VMware vRealize Operations Management Pack for Microsoft Azure**.
- 3 Select either of the following options:
 - To view the objects by region, click **Azure Resources By Region**.
 - To view the objects by cloud account, click **Azure Resources By Subscription**.
- 4 To view the object information by region, region per cloud account, subregion, cloud account, or resource group, select either of the following options:
 - If you are viewing objects by region, select a region. You can click the **Azure Region per Subscription** tab to view the object information for the region per cloud account. You can also expand the inventory tree for each region and select a subregion.
 - If you are viewing objects by cloud account, select a cloud account. You can also expand the inventory tree for each cloud account and select a resource group.

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

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

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

Application Monitoring

You can monitor application services in vRealize Operations Manager. You can also manage the life cycle of agents and application services on virtual machines.

For example, as an administrator, you might need to ensure that the infrastructure provided for running the application services is sufficient and that there are no problems. If you receive a complaint that a particular application service is not working properly or is slow, you can troubleshoot by looking at the infrastructure on which the application is deployed. You can view important metrics related to the applications and share the information with the team managing the applications. You can use vRealize Operations Manager to deploy the agents and send the related application data to vRealize Operations Manager. You can view the data in vRealize Operations Manager and share it with the team so that they can troubleshoot the application service.

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

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

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

Introduction

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

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

The following 23 application services are supported.

Table 4-19.

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

Supported Platforms

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

Platforms Supported by vRealize Operations Manager for Application Monitoring

Platform	Version	Architecture	Application
Red Hat Enterprise Linux	7.x 8.x	64-bit	OS Metrics and all supported applications.
CentOS	7.x	64-bit	OS Metrics and all supported applications.
Windows	Windows Server 2019 Windows Server 2016 Windows 2012 Windows Server 2012 R2	64-bit	OS Metrics and all supported applications.
SUSE Linux Enterprise Server	12.x 15.x	64-bit	OS Metrics and all supported applications.
Oracle Linux	7.x 8.x	64-bit	OS Metrics and all supported applications.
Ubuntu	18.04 LTS 16.04 LTS	64-bit	OS Metrics and all supported applications.
VMware Photon Linux	1.0 2.0 3.0	64-bit	Only OS metrics monitoring supported vRealize Application Remote Collector 8.3 runs on Photon 1.0. vRealize Application Remote Collector 8.2 runs on Photon 1.0. vRealize Application Remote Collector 8.1 runs on Photon 1.0 and vRealize Application Remote Collector 7.5 runs on Photon 1.0 Site Recovery Manager 8.2 runs on Photon 2.0 vSphere- vSphere 6.7 & 6.5 runs on Photon OS 1.0 VMware vSAN 6.7 & VMware vSAN 6.5 runs on Photon OS 1.0 Unified Access Gateway 3.7 runs on Photon 3.0 & 3.6 runs on Photon 2.0.

Sizing Reference Data

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

For more information, see the Knowledge Base article [2093783](#).

Supported Versions of Application Services

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

Application Versions Validated to Work for Application Monitoring

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

Application Name	Versions Validated in the Lab
Redis	5:4.0.9-1ubuntu0.2
Riak	2.1.4 2.2.3
SharePoint	2013
Apache Tomcat	9.0.17 9.0.22 8.0.33 7.0.92
Weblogic	12.2.1.3.0
Websphere	9.0 8.5.5
NTP	4.2.8p10 4.2.6p5
Active Directory	2016 2019
Hyper-V	10.0.17763.1
Cassandra Database	3.11.6 3.11.7
Oracle Database	12c 11c
Velocloud	4.0.0

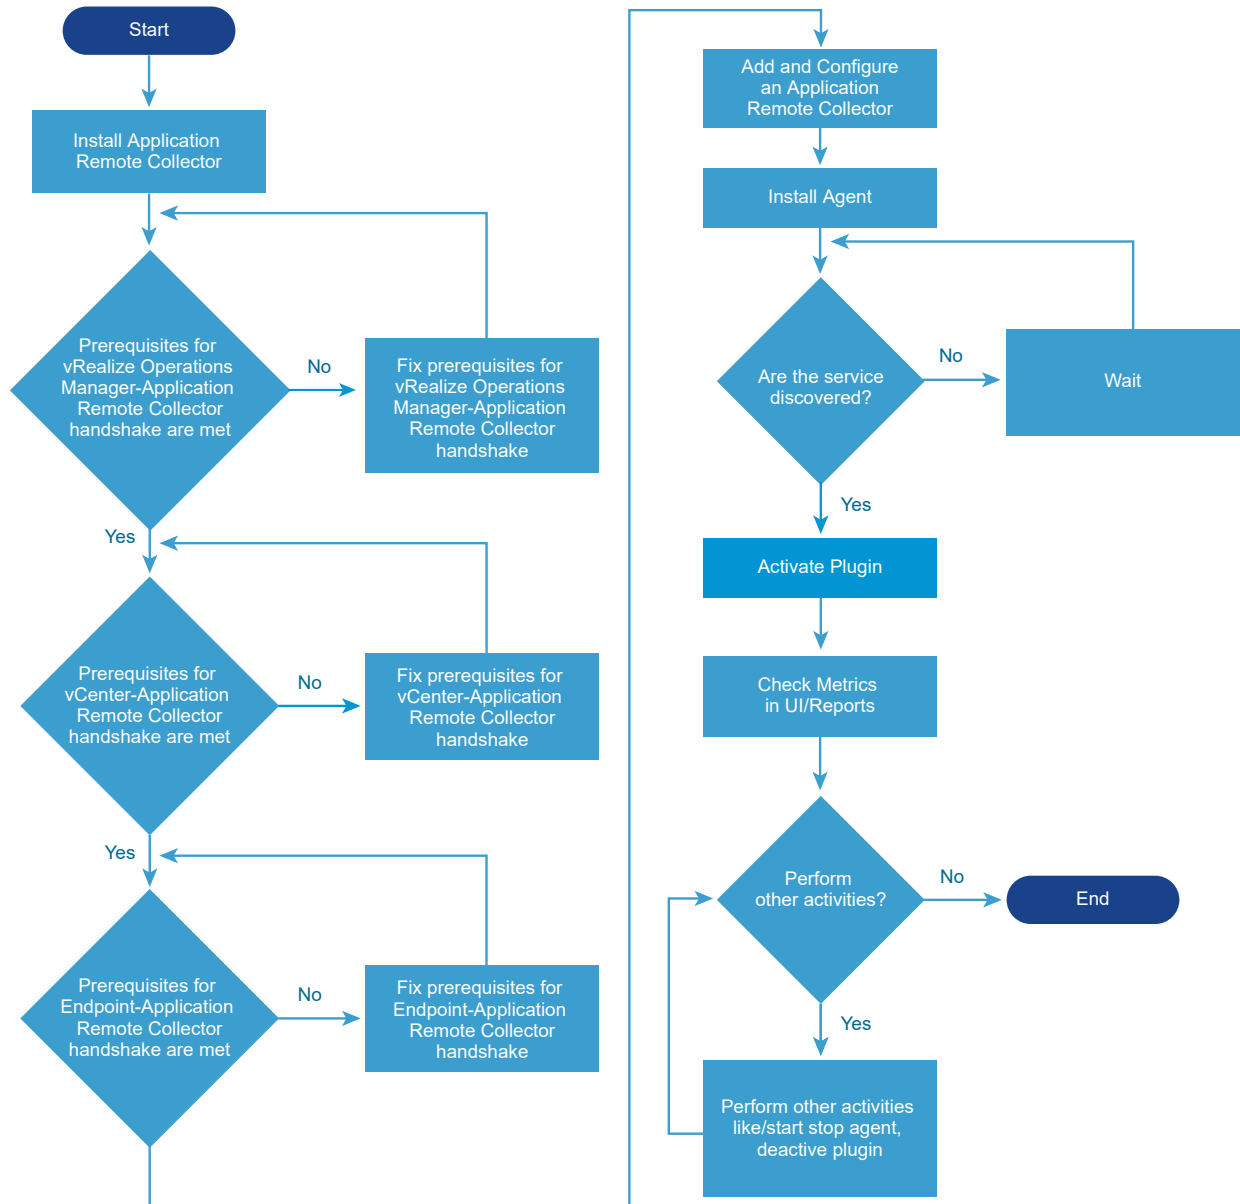
Supported Versions of vCenter Server and VMware Cloud on AWS

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

Steps to Monitor Applications

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

The following flowchart describes how you can set up vRealize Application Remote Collector and vRealize Operations Manager for application monitoring.



Follow these steps to monitor applications.

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

For information about deploying vRealize Application Remote Collector, see [Deploy vRealize Application Remote Collector](#).

- 2 Complete all the prerequisites.

For more, see [Prerequisites](#).

- 3 Add and configure an application remote collector.

For information about configuring vRealize Application Remote Collector, see [Application Remote Collector Page](#) and [Add and Configure an Application Remote Collector](#).

- 4 Install agents on selected VMs.

For more information, see [Install an Agent from the UI](#).

- 5 Activate an application service.

For more information, see [Activate an Application Service](#).

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

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

Deploy or Upgrade vRealize Application Remote Collector

Deploy vRealize Application Remote Collector

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

Prerequisites

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

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

For critical time sourcing, use the Network Time Protocol (NTP). You must ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts, and vRealize Operations Manager.

Procedure

- 1 Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select **Deploy OVF Template**.

The **Deploy OVF Template** wizard opens.

- 2 On the **Deploy OVF template** page, do one of the following and click **Next**:
 - ◆ If you have a URL to the OVA template which is located on the Internet, type the URL in the URL field. Supported URL sources are HTTP and HTTPS.
 - ◆ If you have downloaded the vRealize Application Remote Collector OVA file, click **Local file** and browse to the location of the file and select it.

- 3 On the **Select a name and folder** page, enter a unique name for the virtual machine or vAPP, select a deployment location, and click **Next**.

The default name for the virtual machine is the same as the name of the selected OVF or OVA template. If you change the default name, choose a name that is unique within each vCenter Server virtual machine folder.

The default deployment location for the virtual machine is the inventory object where you started the wizard.

- 4 On the **Select a resource** page, select a resource where to run the deployed VM template, and click **Next**.
- 5 On the **Review details** page, verify the OVF or OVA template details and click **Next**.

Option	Description
Product	vRealize Application Remote Collector.
Version	Version number of the vRealize Application Remote Collector.
Vendor	VMWare.
Publisher	Publisher of the OVF or OVA template, if a certificate included in the OVF or OVA template file specifies a publisher.
Download size	Size of the OVF or OVA file.
Size on disk	Size on disk after you deploy the OVF or OVA template.

- 6 On the **Accept license agreements** page, click **Accept** and then **Next**.
- 7 In the **Select configuration** page, select the size of the deployment.
- 8 On the **Select storage** page, define where and how to store the files for the deployed OVF or OVA template.
 - a Select a VM Storage Policy.
This option is available only if storage policies are enabled on the destination resource.
 - b (Optional) Enable the **Show datastores from Storage DRS clusters** check box to choose individual datastores from Storage DRS clusters for the initial placement of the virtual machine.
 - c Select a datastore to store the deployed OVF or OVA template.
The configuration file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine or vApp and all associated virtual disk files.
- 9 On the **Select networks** page, select a source network and map it to a destination network. Click **Next**. The source network must have a static FQDN name or static DNS.

The Source Network column lists all networks that are defined in the OVF or OVA template.

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

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

- 11 On the **Ready to complete** page, review the page and click **Finish**.
- 12 After the OVA deployment is complete, you can log in to the virtual appliance from vCenter Server. Right click the virtual appliance that you installed. Click **Open Console**. Use the following credentials to log in:

Log In Details	Value
Username	root
Password	vmware

- 13 Change the root user password.

Note To reset the root user password, see the KB article: [2001476](#)

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

What to do next

- Configure NTP Settings.
- Ensure the prerequisites for handshake with vRealize Operations Manager and vCenter Server are met.
- Log in to vRealize Operations Manager and configure application monitoring.

Configure Network Time Protocol Settings

After you install or upgrade to the latest version of vRealize Application Remote Collector, you must set up accurate timekeeping as part of the deployment. If the time settings between vRealize Application Remote Collector and vRealize Operations Manager are not synchronized, you face agent installation and metric collection issues. Ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts, and vRealize Operations Manager using the Network Time Protocol (NTP).

Procedure

- 1 Log in to the vRealize Application Remote Collector appliance and modify the `ntp.conf` file available in `/etc/ntp.conf` by adding following in the following format:

```
server time.vmware.com
```

Note Replace `time.vmware.com` with a suitable time server setting. You can use the FQDN or IP of the time server.

- 2 Enter the following command to start the NTP daemon:

```
systemctl start ntpd
```

- 3 Enter the following command to enable the NTP daemon:

```
systemctl enable ntpd
```

- 4 Run the following command to verify if NTP is configured correctly:

```
ntpstat
```

If NTP is synchronized correctly, you see a message similar to the following:

```
synchronised to NTP server (10.113.60.176) at stratum 3

time correct to within 50 ms

polling server every 64 s
```

Upgrade vRealize Application Remote Collector

Follow the recommended upgrade flow if you have a version of vRealize Operations Manager prior to version 8.3, and version 8.1 of vRealize Application Remote Collector installed. Version 8.1 of vRealize Application Remote Collector is compatible with version 8.1 of vRealize Operations Manager only. Prepare for downtime during the vRealize Application Remote Collector upgrade process. There will be no flow of metrics from the VMs until the upgrade process finishes. After you upgrade vRealize Application Remote Collector, you must update the agents in the endpoints.

Recommended Upgrade Flow

- Upgrade vRealize Operations Manager from version 8.1 to version 8.3.
- Upgrade vRealize Application Remote Collector to version 8.3.

Upgrade an Existing Installation From the VAMI Portal

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

Prerequisites

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

Procedure

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

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

- 2 Click the **Update** tab.
- 3 Click the **Status** tab, click **Actions > Check Updates**.
- 4 Click **Install Updates**.
- 5 After the updates have installed, click **Reboot** in the **System** tab.

Results

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

What to do next

- Update the endpoint agents to discover new services. For more information, see [Additional Operations from the Manage Agents Tab](#).
- To access the virtual machine appliance through ssh, start the sshd service.
- Perform the post-installation tasks.

Upgrade an Existing Installation from an ISO File

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

Prerequisites

- Download the vRealize Application Remote Collector upgrade ISO file called **VMware vRealize Application Remote Collector 8.3.0 (ISO)** from the [official VMware download location](#).
- You must have the root credentials to log in to the VAMI portal before you perform the upgrade.

Procedure

- 1 Upload the vRealize Application Remote Collector upgrade ISO file to the datastore where the vRealize Application Remote Collector appliance is deployed.
- 2 Power off the vRealize Application Remote Collector virtual machine.
- 3 Mount the vRealize Application Remote Collector upgrade ISO file to the virtual machine.
- 4 Power on the vRealize Application Remote Collector virtual machine.

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

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

- 6 Click **Install Updates** under **Status > Updates**.
- 7 After the updates have installed, click **Reboot** in the **System** tab.

Results

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

What to do next

- Update the endpoint agents to discover new services. For more information, see [Additional Operations from the Manage Agents Tab](#).
- To access the virtual machine appliance through ssh, start the sshd service.
- Perform the post-installation tasks.

Prerequisites

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

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

Figure 4-1. Port Information and Communication with vRealize Operations Manager, vCenter Server, and the End Points (Agent Install from the UI)

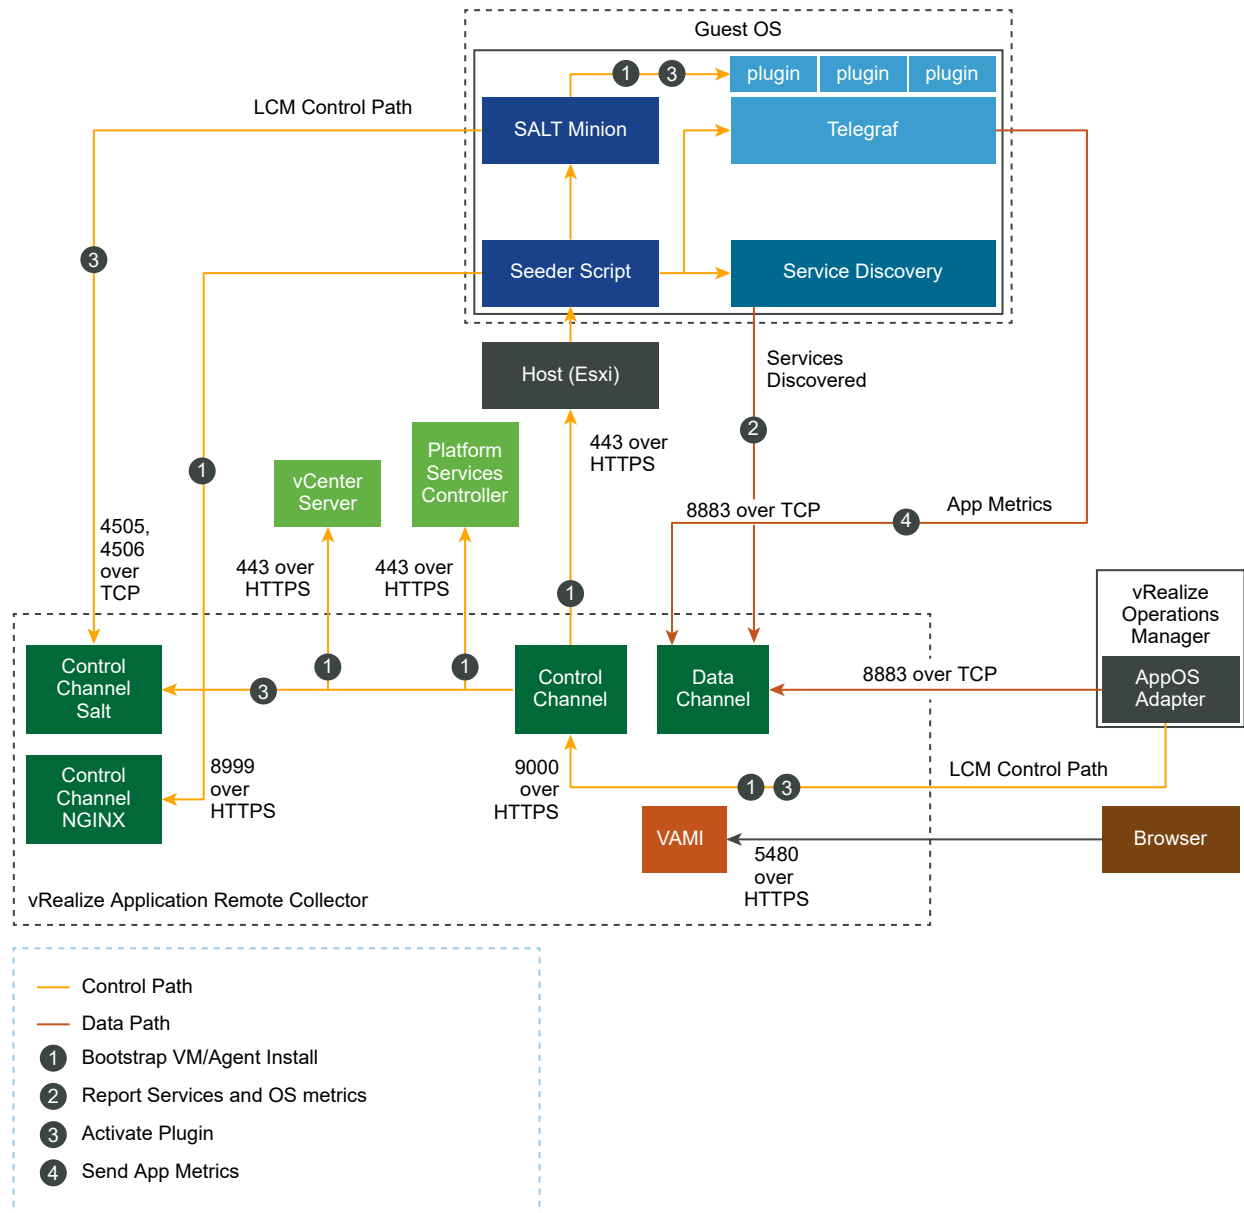
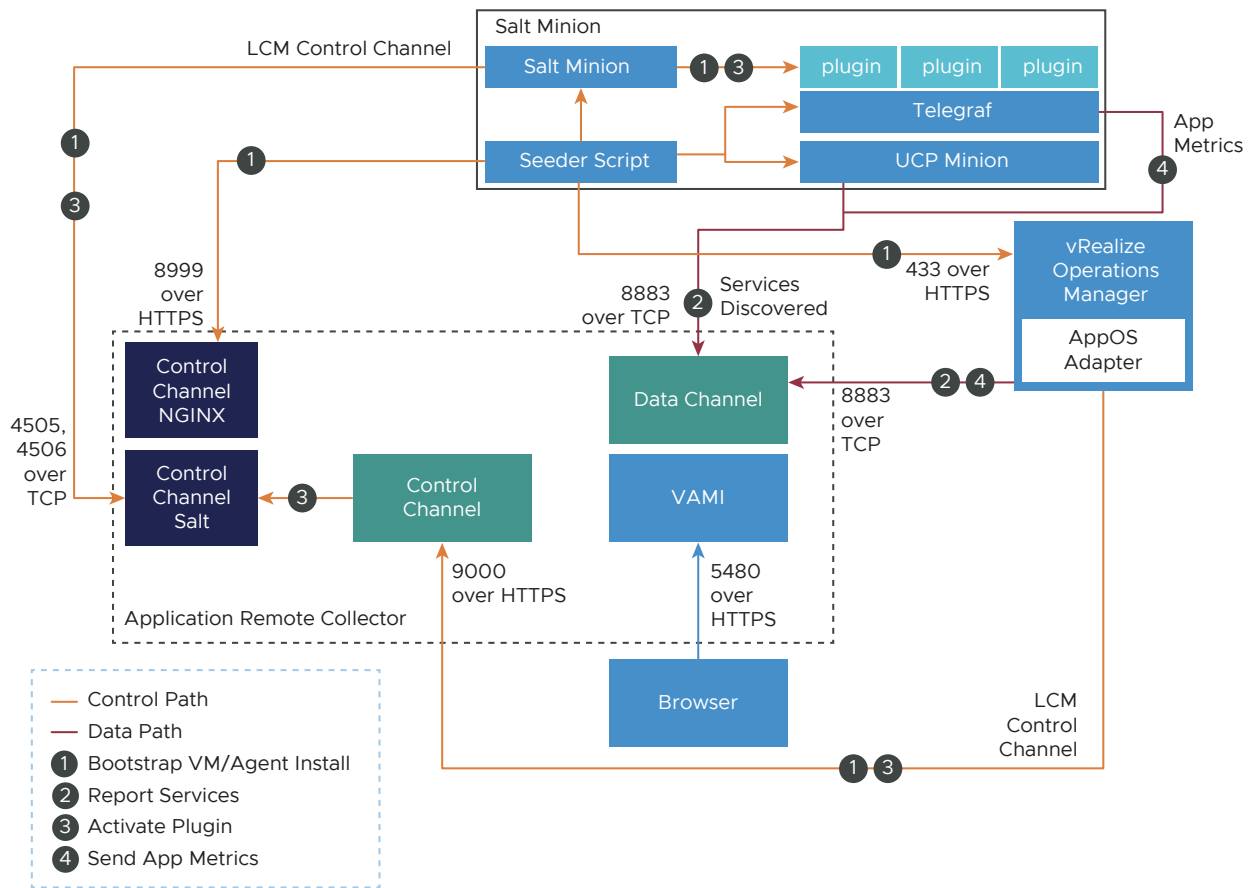


Figure 4-2. Port Information and Communication with the End Points for Script-Based Agent Install



Prerequisites for Communication with vRealize Operations Manager

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

Here are the prerequisites:

- Verify that you have configured a vCenter adapter. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`. See [Install an Agent from the UI](#).
- Ensure that the ports 9000 and 8883 on vRealize Application Remote Collector are reachable from vRealize Operations Manager.
- Download and deploy vRealize Application Remote Collector.

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

For information about deploying the vRealize Application Remote Collector, see [Deploy vRealize Application Remote Collector](#).

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

Prerequisites for Communication with vCenter Server

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

- Ensure that the NTP settings of the ESXi instance that hosts the end points and vRealize Application Remote Collector are in sync.
- Port 443 in vCenter Server is accessible to vRealize Application Remote Collector.
- Port 443 in the ESXi where the workload end-points are deployed must be accessible to vRealize Application Remote Collector.
- Port 443 in Platform Services Controller is accessible to vRealize Application Remote Collector. Open this port if vCenter Server is configured with an external Platform Services Controller.
- Verify that you have configured a vCenter adapter. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have read access at the vCenter Server level and should also have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`. See [Install an Agent from the UI](#).

Prerequisites for Communication with the End Points

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

Here are the prerequisites:

- Ensure that the NTP settings of the ESXi instance that hosts the end points, the end points, and vRealize Application Remote Collector are in sync.
- Ensure that the end points have access to ports 8999, 4505, 4506, and 8883 on vRealize Application Remote Collector.
- Guest operation privileges are required to install agents on virtual machines. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.
- Account privilege prerequisites. See [User Account Prerequisites](#) for more details.
- End-point VM configuration requirements.

- Linux requirements

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

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

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

- Windows 2012 R2 requirement

The end point must be updated with the Universal C Runtime. Refer to the following [link](#) for more information.

- Windows requirement
 - The Visual C++ version must be higher than 14.
 - Performance Monitors on a Windows OS VM must be enabled.
- VMware Tools must be installed and running on the VM on which you want to install the agent. For information about supported VMware Tools versions, click this [link](#).
- If plugin activation requires the location of a file (for example, client certificates for SSL Trust) on the endpoint VM, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

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

User Account Prerequisites

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

Prerequisites for Windows End Points

- To install agents,
 - The user must be either an administrator, or
 - A non-administrator who belongs to the administrator group.

Prerequisites for Linux End Points

- /tmp mount point should be mounted with exec mount option.
- Ensure that the following lines exist in `/etc/sudoers`.

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

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

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

Install User Prerequisites

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

- root user - All privileges
- A non-root user with all privileges -

Password-less sudo elevation access for a non-root user or a non-root user group.

To enable password-less sudo elevation access for a user called *bob*, add *bob*

```
ALL=(ALL:ALL) NOPASSWD: ALL to /etc/sudoers.
```

To enable password-less sudo elevation access for a user group called *bobg*, add *%bobg*

```
ALL=(ALL:ALL) NOPASSWD: ALL to /etc/sudoers.
```

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

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

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

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

Run-Time User Prerequisites

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

- Manually created *arcuser* and *arcgroup*.

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

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

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

```
groupadd arcgroup
```

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

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

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

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

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

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

All privileges:

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

Specific set of privileges:

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

Configure Network Time Protocol Settings

After you install or upgrade to the latest version of vRealize Application Remote Collector, you must set up accurate timekeeping as part of the deployment. If the time settings between vRealize Application Remote Collector and vRealize Operations Manager are not synchronized, you face agent installation and metric collection issues. Ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts, and vRealize Operations Manager using the Network Time Protocol (NTP).

Procedure

- 1 Log in to the vRealize Application Remote Collector appliance and modify the *ntp.conf* file available in */etc/ntp.conf* by adding following in the following format:

```
server time.vmware.com
```

Note Replace *time.vmware.com* with a suitable time server setting. You can use the FQDN or IP of the time server.

- 2 Enter the following command to start the NTP daemon:

```
systemctl start ntpd
```

- 3 Enter the following command to enable the NTP daemon:

```
systemctl enable ntpd
```

- 4 Run the following command to verify if NTP is configured correctly:

```
ntpstat
```

If NTP is synchronized correctly, you see a message similar to the following:

```
synchronised to NTP server (10.113.60.176) at stratum 3  
  
time correct to within 50 ms  
  
polling server every 64 s
```

Add and Configure an Application Remote Collector

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

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

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

For more troubleshooting information on vRealize Application Remote Collector, see [Troubleshooting the Configuration of vRealize Application Remote Collector](#).

Prerequisites

Ensure that you have completed all the prerequisites. For more information, see [Prerequisites](#).

Procedure

- 1 To configure a vRealize Application Remote Collector, click the **Add** icon from the **Application Remote Collector** page.
- 2 In the **Application Remote Collector** page, enter the following details:
 - a FQDN of the vRealize Application Remote Collector you have configured during the installation of vRealize Application Remote Collector.
 - b You cannot modify the user name which is **admin**.
 - c The API password of the vRealize Application Remote Collector you have configured during the installation of vRealize Application Remote Collector.
 - d Click **Next**.

3 From the **Map vCenters** page, complete the following steps:

- a Select the vCenter Servers to which you want to map the vRealize Application Remote Collector.

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

- b The vCenter Servers that are mapped to the vRealize Application Remote Collector are displayed on the page.

- c Click **Test Connection** to validate the connection. The **Review and Accept Certificate** dialog box is displayed. Click **Accept** if you trust the certificate.

If the mapped vCenter Server turns red, it signifies that vRealize Operations Manager cannot communicate with the vRealize Application Remote Collector. If the mapped vCenter Server turns green, it signifies that vRealize Operations Manager can communicate with the vRealize Application Remote Collector.

- d Click **Next**.

4 From the **Summary** page, you view details such as the FQDN, user name, and the vCenter Servers that are mapped to an instance of the vRealize Application Remote Collector.

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

- a Click **Finish**.

What to do next

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

Application Remote Collector Page

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

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

Table 4-20. Options

Options	Description
Add	<p>You can map a vCenter Server with a vRealize Application Remote Collector as part of the configuration process. For more information, see Add and Configure an Application Remote Collector.</p> <p>When you click Test Connection to validate the connection, the Review and Accept Certificate dialog box is displayed. Click Accept if you trust the certificate.</p>
Edit	<p>You can modify the vRealize Application Remote Collector configuration details or the details of the vCenter Servers that are managed.</p> <p>After you modify the details and click Test Connection, the Review and Accept Certificate dialog box is displayed if you have not already accepted the certificate. Click Accept if you trust the certificate. The connection is then validated.</p>
Delete	<p>You can delete the application remote collector. Ensure that you uninstall the agents from the VMs that are monitored before you delete the application remote collector.</p>
Download	<p>You can download vRealize Application Remote Collector. For information about deploying vRealize Application Remote Collector, see Deploy vRealize Application Remote Collector.</p>

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

Table 4-21. Data Grid Options

Option	Description
Name	Displays the FQDN of the vRealize Application Remote Collector.
Application Remote Collector Version	Displays the version of vRealize Application Remote Collector. A gray dot is displayed if there is a newer version of vRealize Application Remote Collector available.
vCenters Managed	Displays the number of vCenter Servers mapped to the vRealize Application Remote Collector.
Collector Server Status	<p>Indicates the health of the vRealize Application Remote Collector.</p> <ul style="list-style-type: none"> ■ Green. Indicates that the vRealize Application Remote Collector is healthy. ■ Red. Indicates that the vRealize Application Remote Collector is not healthy. <p>Point to this cell to view a tooltip that displays the cause if the health status is red.</p> <p>The progress status is displayed when data collection has not started.</p>

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

Install an Agent

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

Install an Agent from the UI

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

Prerequisites

Ensure that you have completed all the prerequisites. For more information, see [Prerequisites](#).

Procedure

- 1 From the **Manage Agents** tab, click the **Install** icon. You see the **Manage Agent** dialog box.
- 2 From the **How do you want to provide VM Credentials** page, complete the following steps:
 - a If you have a common user name and password for all the VMs, select the **Common username and password** option.
 - b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
 - c Click **Next**.
- 3 From the **Provide Credentials** page, depending on whether you have a common credential for all VMs or different credentials for all VMs, enter the following details:
 - a If the selected VMs have a common user name and password, enter the common user name and password.
 - b For different user names and passwords for each VM, download the CSV template and add the required details such as the user name, password for each VM. Use the **Browse** button to select the template.
 - c The **Create run time user on Linux virtual machines, with required permissions as part of agent installation** check box is selected by default. For more information, see [User Account Prerequisites](#).
 - d Click **Next**.
- 4 From the **Summary** page, you can view the list of VMs on which the agent is to be deployed.
- 5 Click **Install Agent**. Refresh the UI to view the agents that are installed.

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

UAC Enabled Machines on Windows End Points

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

- a From `C:\VMware\UCP\downloads`, run a bootstrap launcher.
- b Go to `%SYSTEMDRIVE%\VMware\UCP\downloads`.
- c Open `cmd` with administrator privileges.
- d Run the `cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1` command.
- e View the results from `uaf_bootstrap.log`.
- f Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Agents** tab.

What to do next

You can manage the services on each agent.

For information about uninstalling an agent, see [Uninstall an Agent](#).

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

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

Prerequisites

- Ensure that the end point is available in vRealize Operations Manager.
- Ensure that you have completed all the prerequisites. For more information, see [Prerequisites](#).
- Ensure that the unzip package is available on the VM.
- Ensure that the user has access permissions to the download folder.
- Ensure that the guest IP is properly configured and is unique across vCenter Servers. If more than one VM with the same IP across vCenter Servers is monitored, the script cannot resolve and subscribe to application monitoring.
- Ensure that the cloud account is configured for the vCenter Server to which the VM belongs. The vCenter Server must be mapped with vRealize Application Remote Collector.
- Ensure that port 443 in vRealize Operations Manager is accessible to the end point.
- Only IPv4 is supported at present.

Procedure

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

Run one of the following commands:

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

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

- 2 Make the script executable by running the following command:

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

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

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

Description of arguments:

operation - Bootstrap operation. values: install, uninstall.

vrops_ip_or_fqdn - IP/FQDN of vRealize Operations Manager. This can be the address of any vRealize Operations Manager node or VIP of vRealize Operations Manager.

vrops_user - vRealize Operations Manager user. The user should have enough permissions.

vrops_password - Password of vRealize Operations Manager.

download_tmp_dir - Temporary directory to download agent related bits. It's an optional parameter. Default value: current directory.

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

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

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

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

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

Prerequisites

- Ensure that the end point is available in vRealize Operations Manager.
- Ensure that you have completed all the prerequisites. For more information, see [Prerequisites](#).
- Ensure that the unzip package is available on the VM.
- Ensure that the user has access permissions to the download folder.

- Ensure that the Windows PowerShell is ≥ 4.0 .
- Ensure that the guest IP is properly configured and is unique across vCenter Servers. If more than one VM with the same IP across vCenter Servers is monitored, the script cannot resolve and subscribe to application monitoring.
- Ensure that the cloud account is configured for the vCenter Server to which the VM belongs. The vCenter Server must be mapped with vRealize Application Remote Collector.
- Ensure that port 443 in vRealize Operations Manager is accessible to the end point.
- Only IPv4 is supported at present.

Procedure

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

Run one of the following commands:

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

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

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

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

Description of arguments:

`operation` - Bootstrap operation. values: `install`, `uninstall`.

`vrops_ip_or_fqdn` - IP/FQDN of vRealize Operations Manager. This can be the address of any vRealize Operations Manager node or VIP of vRealize Operations Manager.

`vrops_user` - vRealize Operations Manager user. The user should have enough permissions.

`vrops_password` - Password of vRealize Operations Manager.

`download_tmp_dir` - Temporary directory to download agent related bits. It is an optional parameter. Default value: current directory.

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

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

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

Activate an Application Service

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

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

Prerequisite

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

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

Activate an Application Service

To monitor an application service, complete the following steps:

- 1 Navigate to the **Inventory > Manage Agents** tab.
- 2 Select the VM on which agent is already installed.
- 3 Select **Manage Service** icon and then from the drop-down menu select the **service name**.
- 4 Activate the application service from the right pane of the **Manage <service name> Agent** dialog box.
- 5 Click the **Add** icon in the left pane to add multiple instances of the application service.
- 6 Click the **Delete** icon in the left pane to delete instances of the application service.
- 7 Enter the details for each instance that you add and click **Save**. For configuration details of each application, see [Configuring Supported Application Services](#).

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

The following special characters are permitted in the DB user field: ' [] {} () , . < > ? : ! | / ~ @ # \$ % ^ & * - _ + =

You can provide DB name lists in the following format ['DBNAME_1', 'DBNAME_2', 'DBNAME_3'] where DBNAME_1, DBNAME_2, DBNAME_3 must not contain quotes such as ' and ".

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

For information about deactivating a service, see [Deactivate an Application Service](#).

Configuring Supported Application Services

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

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

Active Directory

Active Directory is supported in vRealize Operations Manager.

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

Active MQ

ActiveMQ is supported in vRealize Operations Manager.

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

Apache HTTPD

Apache HTTPD is supported in vRealize Operations Manager.

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

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

Cassandra Database

Cassandra database is supported in vRealize Operations Manager.

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

Hyper-V

Hyper-V is supported in vRealize Operations Manager.

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

Java

Java is supported in vRealize Operations Manager.

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

JBoss

JBoss is supported in vRealize Operations Manager.

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

MongoDB

MongoDB is supported in vRealize Operations Manager.

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

MS Exchange

MS Exchange is supported in vRealize Operations Manager.

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

MS IIS

MS IIS is supported in vRealize Operations Manager.

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

MS SQL

MS SQL is supported in vRealize Operations Manager.

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

MySQL

MySQL is supported in vRealize Operations Manager.

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

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

NTPD

NTPD is supported in vRealize Operations Manager.

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

Oracle Database

Oracle database is supported in vRealize Operations Manager.

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

Pivotal Server

Pivotal Server is supported in vRealize Operations Manager.

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

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

Postgres

Postgres is supported in vRealize Operations Manager.

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

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

RabbitMQ

RabbitMQ is supported in vRealize Operations Manager.

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

Riak

Riak is supported in vRealize Operations Manager.

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

Sharepoint

Sharepoint is supported in vRealize Operations Manager.

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

Tomcat

Tomcat is supported in vRealize Operations Manager.

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

Weblogic

Weblogic is supported in vRealize Operations Manager.

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

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

Websphere

Websphere is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
IBM Websphere Server URL	Yes	Example : http://localhost:9081
Websphere Authorization Token	Yes	<p>To generate the token, follow the below steps:</p> <ul style="list-style-type: none"> ■ Go to https://www.base64encode.org. ■ Type in the user and password created in the format: user:password ■ Click the Encode button. ■ Copy the resulting Base64 encoded string. Example: d2F2ZWZyb250OndhdmVmcm9udA==

Remote Checks

HTTP Remote Check

HTTP is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display name of the remote check instance.
URL	Yes	http://localhost
Method	Yes	GET/POST/PUT
Proxy	No	Proxy URL: http://localhost

Name	Mandatory?	Comment
Response Timeout	No	Timeout for the connection in seconds. For example, 10.
Follow Redirects	No	True/False if redirects from the server. For example, true/false (all small values).
Body	No	HTTP request body.
Response String Match	No	Substring or regex match in the response body.
SSL CA	No	Path to the SSL CA file on the end point.
SSL Certificates	No	Path to the SSL certificate file on the end point.
SSL Key	No	Path to the SSL key file on the end point.
Skip Host & chain verification	No	Use SSL but skip chain and host verification. Expected: True/False.

ICMP Remote Check

ICMP is supported in vRealize Operations Manager.

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

TCP Remote Check

TCP is supported in vRealize Operations Manager.

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

UDP Remote Check

UDP is supported in vRealize Operations Manager.

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

Configuring Supported VeloCloud Services

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

VeloCloud Orchestrator

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

- Java Application
- VeloCloud Orchestrator
- Nginx

- Clickhouse
- Network Time Protocol
- MySQL
- Redis

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

- Backend
- Portal
- Upload

VeloCloud Orchestrator details.

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

Nginx

Nginx is supported in vRealize Operations Manager.

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

ClickHouse

ClickHouse is supported in vRealize Operations Manager.

Name	Mandatory?	Comment
Display Name	Yes	Display Name of the application instance.
Servers URL	Yes	http://127.0.0.1:8123

Name	Mandatory?	Comment
User name	No	User name for the ClickHouse service.
Password	No	Password

NTPD

NTPD is supported in vRealize Operations Manager.

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

MySQL

MySQL is supported in vRealize Operations Manager.

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

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

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

Name	Mandatory?	Comment
Databases	No	Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated. For example, 'database1','database2','database3'.
TLS Connection	No	Accepted values are true, false, and skip-verify.

Redis

Redis is supported in vRealize Operations Manager.

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

VeloCloud Gateway

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

- Network Time Protocol
- VeloCloud Gateway

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

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

VeloCloud Gateway details.

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

NTPD

NTPD is supported in vRealize Operations Manager.

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

Pre-Requirements for Application Services

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

Apache HTTPD

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

```
<IfModule mod_status.c>

<Location /server-status>

    SetHandler server-status

</Location>

ExtendedStatus On

</IfModule>
```

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

```
systemctl restart httpd
```

Java Plugins

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

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

Nginx

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

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

Restart the Nginx service with the following command:

```
systemctl restart nginx
```

Postgres

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

```
sudo service postgresql restart
```

Cassandra Database

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

- 1 Modify `/etc/default/cassandra`.

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

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

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

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

Oracle Database

To monitor the Oracle database, complete these steps:

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

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

2 Create a User.

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

3 Install Python 3.6 or later.

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

4 Set the PATH of TNS_ADMIN.

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

Note Oracle database cannot be activated on Linux platforms.

Active MQ 5.16 and Later Versions

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

- Navigate to /opt/activemq/apache-activemq-5.16.0/webapps/api/WEB-INF/classes/jolokia-access.xml
- Remove or comment out the following lines:

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

- Restart the Active MQ service.

MS SQL

The user account must have the following permissions to monitor the MS SQL application with Telegraf.

```
USE master;
GO
CREATE LOGIN [telegraf] WITH PASSWORD = N'mystrongpassword';
GO
GRANT VIEW SERVER STATE TO [telegraf];
GO
GRANT VIEW ANY DEFINITION TO [telegraf];
GO
```

Additional Operations from the Manage Agents Tab

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

the agents on the VMs. You can also discover and manage the services on each agent that you install.

Where You Manage the Agents

To manage the agents and application services, in the menu, select **Administration**, and then from the left pane select **Inventory**. From the right pane, click the **Manage Agents** tab.

Table 4-22. Options

Options	Description
Install	Installs the agents on the selected VM. Select the VMs on which you want to install the agent and click the Install icon. For more information, see Install an Agent from the UI .
Uninstall	Uninstalls the agent. Select the VMs on which you want to uninstall the agent and click the Uninstall icon. For more information, see Uninstall an Agent .
Update	Updates agents that are at a lower version. Select the VMs on which you want to update the agent and click the Update icon. After the agents are updated, the last operation status changes to Content Upgrade Success .
Start	If you have temporarily stopped sending metrics to vRealize Operations Manager, you can use this option to start data collection for the application service.
Stop	During a maintenance period, you can temporarily stop sending application service metrics to vRealize Operations Manager. Select the VMs on which you want to stop the agent and click the Stop icon.
Manage Service	You can configure and activate the application services that are discovered on the virtual machines where the agents are installed. For configuration details of each application, see Configuring Supported Application Services .
Manage Service > Remote Check	Allows you to enable remote checks such as ICMP Check, UDP Check, TCP Check, and HTTP Check.
Manage Service > Custom Script	Allows you to run custom scripts in the VM and collect custom data which can be then consumed as a metric. For more information, see Custom Script .
Show Detail	Displays the Summary tab of the selected VM.
All Filters	Filters the VMs based on the name of the VM, the operating system it runs on, the application service discovered, and the power status of the VM.

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

Table 4-23. Data Grid Options

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

Table 4-23. Data Grid Options (continued)

Option	Description
Services Discovered/Configured	<p>List of the supported application services discovered on the VM.</p> <ul style="list-style-type: none"> ■ A red dot against the application service indicates that the application service has been activated but there is a problem with data collection. <p>When there is more than one application service of the same kind, and one of them is activated, but the other is not collecting data, a red dot is still displayed against the application service.</p> <ul style="list-style-type: none"> ■ A gray dot before the application service indicates that the agent requires reactivation. The application service must be reactivated. For reactivation, see Activate an Application Service for more information. ■ A gray pause symbol indicates that the agents have stopped. ■ A green icon against the application service indicates that the application service is activated. <p>You see a blue icon with three horizontal dots if there is a problem with the activation. Click the question mark for more information about the warning. The warning is also displayed in the following locations:</p> <ul style="list-style-type: none"> ■ In the Objects tab for the specific application service. Move your cursor over the green icon in the Collection Status column. ■ For the specific application service, click the Show Details option from the Manage Agents tab. Move your cursor over the green icon in the top panel to view the warning message. ■ If an application service has been deactivated or not activated, you see a gray pause symbol displayed against the application service. ■ After you have added the parameters and activated the application service, the progress status is displayed until data collection starts. <p>Click the colored dots for more information about the application services.</p>
Agent Status	<p>Displays the status of the agent at the end point.</p> <ul style="list-style-type: none"> ■ Blue icon. Indicates that the agent is not installed. ■ Green icon. Indicates that the agent is running. ■ Red icon. Indicates that the agent has stopped. ■ Gray dot. Appears in front of the service and indicates that plugin reactivation is required.

Table 4-23. Data Grid Options (continued)

Option	Description
Last Operation Status	<p>Status of the last operation. The possible values are:</p> <ul style="list-style-type: none"> ■ No Operation ■ Install Success ■ Install Failed ■ Install In Progress ■ Start Success ■ Start Failed ■ Start In Progress ■ Stop Success ■ Stop Failed ■ Stop In Progress ■ Update Success ■ Update Failed ■ Update In Progress ■ Uninstall Success ■ Uninstall Failed ■ Uninstall In Progress ■ Download Success
VM State	<p>Power status of the VMs. The possible values are:</p> <ul style="list-style-type: none"> ■ Powered On ■ Powered Off
ARC	FQDN of the instance of the vRealize Application Remote Collector that you are using.
Agent Version	Version of the agent on the VM. A gray dot is displayed if the VM requires an update.
vCenter Name	Name of the vCenter Adapter instance to which that VM resource belongs.

To manage the agent, follow these steps:

- 1 Install the agent.

For more information, see [Install an Agent from the UI](#).

- 2 Manage the application services on each agent.

For more information, see [Configure Application Services](#).

- 3 Stop and start the agents on the VMs.

- 4 Uninstall the agent.

For more information, see [Uninstall an Agent](#).

5 Update agents that are at a lower version.

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

Custom Script

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

Prerequisites

- All the scripts that you run using the custom script, must output a single integer value. If the output is not a single integer value, an error is displayed in the user interface.
- The custom script uses Telegraf's `exec` plugin to run scripts on a VM's operating system. The scripts are run by the user who installed the Telegraf agent on an operating system. In Linux operating systems, a special user called *arcuser* with specific privileges, is created for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that *arcuser* user. Ensure that the *arcuser* can run the scripts that use the custom script (the *arcuser* must have permissions to run the script). For example, the *arcuser* created automatically by vRealize Application Remote Collector, does not have privileges to run scripts which are stored under the `/root` directory.
- The script must be placed in the `/opt/vmware` folder.

Instance Settings

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

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

Note

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

All Metrics Tab

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

Deactivate an Application Service

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

Prerequisite

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

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

Deactivate an Application Service

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

- 1 Navigate to the **Inventory > Manage Agents** tab.
- 2 Select the VM on which the agent is already installed.
- 3 Select the **Manage Service** icon and then from the drop-down menu select the **service name**.

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

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

For information on activating an application service, see [Activate an Application Service](#).

Uninstall an Agent

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

Prerequisites

- Time synchronization between vRealize Application Remote Collector, vRealize Operations Manager, ESX hosts, and Windows and Linux target VMs is mandatory for secure communication.
- Guest operation privileges are required to install agents on virtual machines. The vCenter Server user account with which the vCenter adapter is configured in vRealize Operations Manager, should have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.
- Account privilege prerequisites. See [User Account Prerequisites](#) for more details.
- End-point VM configuration requirements.
 - Linux requirements

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

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

Configure mount point on `/tmp` directory to allow script execution.
 - Windows 2012 R2 requirement

The end point must be updated with the Universal C Runtime. Refer to the following [link](#) for more information.
 - Windows requirement

The Visual C++ version must be higher than 14.
- VMware Tools must be installed and running on the VM on which you want to install the agent.

Procedure

- 1 From the **Manage Agents** tab, click the **Uninstall** icon. You see the **Manage Agent** dialog box.

- 2 From the **How do you want to provide VM Credentials** page, complete the following steps:
 - a If you have a common user name and password for all the VMs, select the **Common username and password** option.
 - b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
 - c Click **Next**.
- 3 From the **Provide Credentials** page, depending on whether you have a common credential for all VMs or different credentials for all VMs, enter the following details:
 - a If your VM has a single user name and password, enter the common user name and password.
 - b For multiple user names and passwords for each VM, download the CSV template and add the details. Use the **Browse** button to select the template.
 - c Click **Next**.
- 4 From the **Summary** page, you can view the list of VMs on which the agent is deployed.
- 5 Click **Uninstall Agent**. Refresh the UI to view the progress of agent uninstallation.

The **Agent Status** and **Services Discovered** columns in the workspace indicate that uninstallation is complete and that there are no application services discovered on each agent.

UAC Enabled Machines on Windows End Points

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

- a From C:\VMware\UCP\downloads, run a bootstrap launcher.
- b Go to %SYSTEMDRIVE%\VMware\UCP\downloads.
- c Open cmd with administrator privileges.
- d Run the cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1 command.
- e View the results from uaf_bootstrap.log.
- f Verify the status of agent uninstallation from the **Agent Status** and **Last Operation Status** columns in the **Manage Agents**

For information about installing an agent, see [Install an Agent from the UI](#).

Configure Application Services

You can configure the application services on the VMs where the agents are installed.

Procedure

- 1 Select a VM on which the agent has been installed and the application services have been discovered, from the **Manage Agents** tab.

- 2 Select **Manage Service** and then from the drop-down menu select the **service name**. You see the **Manage <service name> Agent** dialog box.
- 3 By default, all metrics are collected for the activated application service.
- 4 Activate data collection for the application service.
- 5 Enter the relevant settings for the application service. For configuration details of each application, see [Configuring Supported Application Services](#).
- 6 Click **Save** and then **Close**.

Fields with a star are mandatory.

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

What to do next

You can monitor the applications services from vRealize Operations Manager.

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

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

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

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

Prerequisites

- vRealize Application Remote Collector appliance must be configured with a static I.P. or static FQDN. The endpoints must be configured.
- Back up the network configuration details of the vRealize Application Remote Collector appliance. Capture the network configuration details of vRealize Application Remote Collector either using the VAMI UI or vCenter Server Tools. Keep the network details available when you restore the vRealize Application Remote Collector appliance from the backup.
- The sizing of the new vRealize Application Remote Collector appliance that you are restoring a backup to, should be greater or equal to the old appliance. The network configuration, static I.P. or static FQDN should be the same. This is to enable the endpoint VMs to reach the new appliance.

Procedure

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

- a Connect to the virtual machine running vRealize Application Remote Collector using SSH.
- b Enter the following command to access the scripts folder:

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

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

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

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

- d Archive the `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file to a remote location.
- 2 A cron job also runs every day and backs up the `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file. The `.tar` file is stored for five days. On the sixth day, the oldest `.tar` file is deleted and replaced. In order to restore the vRealize Application Remote Collector appliance from the `.tar` file, archive the file to a remote location.
 - 3 Restore the backed up configuration files to a new vRealize Application Remote Collector appliance.
 - a Configure the new vRealize Application Remote Collector appliance with the same network and IP configuration as the previous appliance. This information is available in the network configuration file that you backed up.
 - b Connect to the VM running vRealize Application Remote Collector using SSH.
 - c Retrieve the latest `Application-Remote-Collector-State-Bundle_<<Timestamp>>.tar` file from the archive, and copy it to a location which is accessible by the vRealize Application Remote Collector appliance.

- d Enter the following command to access the scripts folder:

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

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

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

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

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

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

Results

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

What to do next

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

Summary of Discovered and Supported Operating Systems and Application Services

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

Where You View Applications in vRealize Operations Manager

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

Discovered Operating Systems and Services

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

Supported Operating Systems

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

Supported Services

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

Metrics Collected

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

Operating System Metrics

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

Linux Platforms

The following metrics are collected for Linux operating systems:

Table 4-24. Metrics for Linux

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

Table 4-24. Metrics for Linux (continued)

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

Table 4-24. Metrics for Linux (continued)

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

Windows Platforms

The following metrics are collected for Windows operating systems:

Table 4-25. Metrics for Windows

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

Table 4-25. Metrics for Windows (continued)

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

Application Service Metrics

Metrics are collected for 20 application services.

Active Directory Metrics

Metrics are collected for the Active Directory application service.

Table 4-26. Active Directory Metrics

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

Table 4-26. Active Directory Metrics (continued)

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

No metrics are collected for the category Active Directory.

Apache Tomcat

Metrics are collected for the Apache Tomcat application service.

Table 4-27. Apache Tomcat

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

Table 4-27. Apache Tomcat (continued)

Metric Name	Category	KPI
JVM Memory Pool:<InstanceName> Peak Usage Committed Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Initial Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Maximum Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Used Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Committed Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Initial Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Maximum Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Used Memory	Tomcat Server	False
Process CPU Usage (%)	Tomcat Server	True
System CPU Usage (%)	Tomcat Server	True
System Load Average (%)	Tomcat Server	True
Threading Thread Count	Tomcat Server	False
Uptime	Tomcat Server	True
JSP Count	Tomcat Server Web Module	False
JSP Reload Count	Tomcat Server Web Module	False
JSP Unload Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Error Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Processing Time	Tomcat Server Web Module	False
Cache : Hit Count	Tomcat Server Web Module	False
Cache : Lookup Count	Tomcat Server Web Module	False
Current Thread Count	Tomcat Server Global Request Processor	True
Current Threads Busy	Tomcat Server Global Request Processor	True

Table 4-27. Apache Tomcat (continued)

Metric Name	Category	KPI
errorRate	Tomcat Server Global Request Processor	False
Total Request Bytes Received	Tomcat Server Global Request Processor	False
Total Request Bytes Sent	Tomcat Server Global Request Processor	False
Total Request Count	Tomcat Server Global Request Processor	True
Total Request Error Count	Tomcat Server Global Request Processor	True
Total Request Processing Time	Tomcat Server Global Request Processor	False

MS SQL Metrics

Metrics are collected for the MS SQL application service.

Table 4-28. MS SQL Metrics

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

Table 4-28. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance Buffer Manager Page lookups per second	Microsoft SQL Server	False
Performance Buffer Manager Page reads per second	Microsoft SQL Server	False
Performance Buffer Manager Page writes per second	Microsoft SQL Server	False
Performance Databases Active Transactions	Microsoft SQL Server	True
Performance Databases Data File(s) Size	Microsoft SQL Server	True
Performance Databases Log Bytes Flushed/Sec	Microsoft SQL Server	False
Performance Databases Log File(s) Size	Microsoft SQL Server	False
Performance Databases Log File(s) Used Size	Microsoft SQL Server	False
Performance Databases Log Flush Wait Time	Microsoft SQL Server	False
Performance Databases Log Flushes per second	Microsoft SQL Server	False
Performance Databases Transactions per second	Microsoft SQL Server	False
Performance Databases Write Transactions per second	Microsoft SQL Server	False
Performance Databases XTP Memory Used	Microsoft SQL Server	False
Performance General Statistics Active temp Tables	Microsoft SQL Server	False
Performance General Statistics Logins per second	Microsoft SQL Server	False
Performance General Statistics Logouts per second	Microsoft SQL Server	False
Performance General Statistics Processes Blocked	Microsoft SQL Server	False
Performance General Statistics Temp Tables Creation Rate	Microsoft SQL Server	False
Performance General Statistics User Connections	Microsoft SQL Server	False
Performance Locks Average Wait Time	Microsoft SQL Server	False

Table 4-28. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance Locks Lock Requests per second	Microsoft SQL Server	False
Performance Locks Lock Wait Time	Microsoft SQL Server	True
Performance Locks Lock Waits per second	Microsoft SQL Server	True
Performance Locks Number of Deadlocks per second	Microsoft SQL Server	True
Performance Memory Manager Connection Memory	Microsoft SQL Server	False
Performance Memory Manager Lock Memory	Microsoft SQL Server	False
Performance Memory Manager Log Pool Memory	Microsoft SQL Server	False
Performance Memory Manager Memory Grants Pending	Microsoft SQL Server	True
Performance Memory Manager SQL Cache Memory	Microsoft SQL Server	False
Performance Memory Manager Target Server Memory	Microsoft SQL Server	True
Performance Memory Manager Total Server Memory	Microsoft SQL Server	True
Performance Resource Pool Stats internal Active memory grant amount	Microsoft SQL Server	False
Performance Resource Pool Stats internal CPU Usage Percentage (%)	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read Bytes per second	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read IO	Microsoft SQL Server	False
Wait Stats:<InstanceName> Wait Time (ms)	Microsoft SQL Server	False
Wait Stats<InstanceName> Number of Waiting tasks (ms)	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read IO Throttled Per Second	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Write Bytes per second (Bps)	Microsoft SQL Server	False

Table 4-28. MS SQL Metrics (continued)

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

There are no metrics collected for Microsoft SQL Server Database.

PostgreSQL

Metrics are collected for the PostgreSQL application service.

Table 4-29. PostgreSQL

Metric Name	Category	KPI
Buffers Buffers Allocated	PostgreSQL	False
Buffers Buffers Written by Backend	PostgreSQL	True
Buffers Buffers Written by Background Writer	PostgreSQL	True

Table 4-29. PostgreSQL (continued)

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

IIS Metrics

Metrics are collected for the IIS application service.

Table 4-30. IIS Metrics

Metric Name	Category	KPI
HTTP Service Request Queues<InstanceName>AppPool CurrentQueueSize	IIS HTTP Service Request Queues	True
HTTP Service Request Queues<InstanceName>AppPool RejectedRequests	IIS HTTP Service Request Queues	False
Web Services<InstanceName> Web Site Bytes Received	IIS Web Services	False
Web Services<InstanceName> Web Site Bytes Sent/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Bytes Total/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Connection Attempts/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Current Connections	IIS Web Services	False
Web Services<InstanceName> Web Site Get Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Locked Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Not Found Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Post Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Service Uptime	IIS Web Services	False
Web Services<InstanceName> Web Site Total Bytes Sent	IIS Web Services	False
Web Services<InstanceName> Web Site Total Get Requests	IIS Web Services	True
Web Services<InstanceName> Web Site Total Post Requests	IIS Web Services	True
Web Services<InstanceName> Web Site Total Put Requests	IIS Web Services	False
Current File Cache Memory Usage (bytes)	IIS Web Services Cache	False
File Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Misses	IIS Web Services Cache	False
Total Flushed URIs	IIS Web Services Cache	False

Table 4-30. IIS Metrics (continued)

Metric Name	Category	KPI
URI Cache Hits	IIS Web Services Cache	False
URI Cache Hits Percent (%)	IIS Web Services Cache	False
URI Cache Misses	IIS Web Services Cache	False
ASP.NET<InstanceName> Application Restarts	IIS ASP.NET	True
ASP.NET<InstanceName> Request Wait Time	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Current	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Queued	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Rejected	IIS ASP.NET	True
MS.NET<InstanceName> Allocated Bytes/sec	MS.NET	True
MS.NET<InstanceName> Current Queue Length	MS.NET	False
MS.NET<InstanceName> Finalization Survivors	MS.NET	False
MS.NET<InstanceName> Gen 0 Collections	MS.NET	False
MS.NET<InstanceName> Gen 0 heap size	MS.NET	False
MS.NET<InstanceName> Gen 1 Collections	MS.NET	False
MS.NET<InstanceName> Gen 1 heap size	MS.NET	False
MS.NET<InstanceName> Gen 2 Collections	MS.NET	False
MS.NET<InstanceName> Gen 2 heap size	MS.NET	False
MS.NET<InstanceName> IL Bytes Jitted / sec	MS.NET	False
MS.NET<InstanceName> Induced GC	MS.NET	False
MS.NET<InstanceName> Large Object Heap size	MS.NET	False
MS.NET<InstanceName> No of current logical Threads	MS.NET	True

Table 4-30. IIS Metrics (continued)

Metric Name	Category	KPI
MS.NET<InstanceName> No of current physical Threads	MS.NET	True
MS.NET<InstanceName> No of current recognized threads	MS.NET	False
MS.NET<InstanceName> No of Exceps Thrown / sec	MS.NET	True
MS.NET<InstanceName> No of total recognized threads	MS.NET	False
MS.NET<InstanceName> Percent Time in Jit	MS.NET	False
MS.NET<InstanceName> Pinned Objects	MS.NET	False
MS.NET<InstanceName> Stack Walk Depth	MS.NET	False
MS.NET<InstanceName> Time in RT checks	MS.NET	False
MS.NET<InstanceName> Time Loading	MS.NET	True
MS.NET<InstanceName> Total No of Contentions	MS.NET	False
MS.NET<InstanceName> Total Runtime Checks	MS.NET	True

MS Exchange Server Metrics

Metrics are collected for the MS Exchange Server application service.

Table 4-31. MS Exchange Server Metrics

Metric Name	Category	KPI
Active Manager Server Active Manager Role	MS Exchange	False
Active Manager Server Database State Info Writes per second	MS Exchange	False
Active Manager Server GetServerForDatabase Server-Side Calls	MS Exchange	False
Active Manager Server Server-Side Calls per second	MS Exchange	True
Active Manager Server Total Number of Databases	MS Exchange	True
ActiveSync Average Request Time	MS Exchange	True
ActiveSync Current Requests	MS Exchange	False

Table 4-31. MS Exchange Server Metrics (continued)

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

Table 4-31. MS Exchange Server Metrics (continued)

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

JBoss EAP Metrics

Metrics are collected for the JBoss EAP application service.

Table 4-32. JBoss EAP Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Jboss Server	False
Buffer Pool<InstanceName> Memory Used	Jboss Server	False
Buffer Pool<InstanceName> Total Capacity	Jboss Server	False
Class Loading Loaded Class Count	Jboss Server	False
Class Loading Total Loaded Class Count	Jboss Server	False
Class Loading Unloaded Class Count	Jboss Server	False

Table 4-32. JBoss EAP Metrics (continued)

Metric Name	Category	KPI
File Descriptor Usage Max File Descriptor Count	Jboss Server	False
File Descriptor Usage Open File Descriptor Count	Jboss Server	False
Http Listener<InstanceName> Bytes Received	Jboss Server	False
Http Listener<InstanceName> Bytes Sent	Jboss Server	False
Http Listener<InstanceName> Error Count	Jboss Server	False
Http Listener<InstanceName> Request Count	Jboss Server	False
Https Listener<InstanceName> Bytes Received	Jboss Server	False
Https Listener<InstanceName> Bytes Sent	Jboss Server	False
Https Listener<InstanceName> Error Count	Jboss Server	False
Https Listener<InstanceName> Request Count	Jboss Server	False
Process CPU Usage (%)	Jboss Server	False
System CPU Usage (%)	Jboss Server	False
System Load Average (%)	Jboss Server	False
Threading Daemon Thread Count	Jboss Server	False
Threading Peak Thread Count	Jboss Server	False
Threading Thread Count	Jboss Server	False
Threading Total Started Thread Count	Jboss Server	False
Uptime	Jboss Server	False
UTILIZATION Heap Memory Usage	Jboss Server	False
Garbage Collection<InstanceName> Total Collection Count	Jboss JVM Garbage Collector	False
Garbage Collection<InstanceName> Total Collection Time	Jboss JVM Garbage Collector	False
JVM Memory Heap Memory Usage Committed Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Initial Memory	Jboss JVM Memory	False

Table 4-32. JBoss EAP Metrics (continued)

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

RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

Table 4-33. RabbitMQ Metrics

Metric Name	Category	KPI
CPU Limit	RabbitMQ	False
CPU Used	RabbitMQ	True
Disk Free	RabbitMQ	False
Disk Free limit	RabbitMQ	False
FileDescriptor Total	RabbitMQ	False
FileDescriptor Used	RabbitMQ	False
Memory Limit	RabbitMQ	False
Memory Used	RabbitMQ	True
Messages Acked	RabbitMQ	False
Messages Delivered	RabbitMQ	False
Messages Delivered get	RabbitMQ	False
Messages Published	RabbitMQ	False
Messages Ready	RabbitMQ	False
Messages Unacked	RabbitMQ	False
Socket Limit	RabbitMQ	False
Socket Used	RabbitMQ	True
UTILIZATION Channels	RabbitMQ	True
UTILIZATION Connections	RabbitMQ	True
UTILIZATION Consumers	RabbitMQ	True
UTILIZATION Exchanges	RabbitMQ	True
UTILIZATION Messages	RabbitMQ	True
UTILIZATION Queues	RabbitMQ	True
Messages Publish in	RabbitMQ Exchange	False
Messages Publish out	RabbitMQ Exchange	False
Consumer Utilisation	RabbitMQ Queue	False
Consumers	RabbitMQ Queue	False

Table 4-33. RabbitMQ Metrics (continued)

Metric Name	Category	KPI
Memory	RabbitMQ Queue	False
Messages Ack	RabbitMQ Queue	False
Messages Ack rate	RabbitMQ Queue	False
Messages Deliver	RabbitMQ Queue	False
Messages Deliver get	RabbitMQ Queue	False
Messages Persist	RabbitMQ Queue	False
Messages Publish	RabbitMQ Queue	False
Messages Publish rate	RabbitMQ Queue	False
Messages Ram	RabbitMQ Queue	False
Messages Ready	RabbitMQ Queue	False
Messages Redeliver	RabbitMQ Queue	False
Messages Redeliver rate	RabbitMQ Queue	False
Messages Space	RabbitMQ Queue	False
Messages Unack	RabbitMQ Queue	False
Messages Unacked	RabbitMQ Queue	False
Messages	RabbitMQ Queue	False

There are no metrics collected for RabbitMQ Virtual Host.

MySQL Metrics

Metrics are collected for the MySQL application service.

Table 4-34. MySQL Metrics

Metric Name	Category	KPI
Aborted connection count	MySQL	True
Connection count	MySQL	True
Event wait average time	MySQL	False
Event wait count	MySQL	False
Binary Files Binary Files Count	MySQL	False
Binary Files Binary Size Bytes	MySQL	False
Global Status Aborted Clients	MySQL	False

Table 4-34. MySQL Metrics (continued)

Metric Name	Category	KPI
Global Status Binlog Cache Disk Use	MySQL	False
Global Status Bytes Received	MySQL	False
Global Status Bytes Sent	MySQL	False
Global Status Connection Errors Accept	MySQL	False
Global Status Connection Errors Internal	MySQL	False
Global Status Connection Errors Max Connections	MySQL	False
Global Status Queries	MySQL	False
Global Status Threads Cached	MySQL	False
Global Status Threads Connected	MySQL	False
Global Status Threads Running	MySQL	False
Global Status Uptime	MySQL	False
Global Variables Delayed Insert Limit	MySQL	False
Global Variables Delayed Insert Timeout	MySQL	False
Global Variables Delayed Queue Size	MySQL	False
Global Variables Max Connect Errors	MySQL	False
Global Variables Max Connections	MySQL	False
Global Variables Max Delayed Threads	MySQL	False
Global Variables Max Error Count	MySQL	False
InnoDB All deadlock count	MySQL	False
InnoDB Buffer Pool Bytes Data	MySQL	False
InnoDB Buffer Pool Bytes Data	MySQL	False
InnoDB Buffer Pool Bytes Dirty	MySQL	False
InnoDB Buffer Pool Dump Status	MySQL	False
InnoDB Buffer Pool Load Status	MySQL	False
InnoDB Buffer Pool Pages Data	MySQL	False
InnoDB Buffer Pool Pages Dirty	MySQL	False

Table 4-34. MySQL Metrics (continued)

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

NGINX Metrics

Metrics are collected for the NGINX application service.

Table 4-35. NGINX Metrics

Metric Name	Category	KPI
HTTP Status Info Accepts	Nginx	True
HTTP Status Info Active connections	Nginx	False

Table 4-35. NGINX Metrics (continued)

Metric Name	Category	KPI
HTTP Status Info Handled	Nginx	True
HTTP Status Info Reading	Nginx	False
HTTP Status Info Requests	Nginx	False
HTTP Status Info Waiting	Nginx	True
HTTP Status Info Writing	Nginx	False

Sharepoint Metrics

Metrics are collected for the Sharepoint application service.

Table 4-36. Sharepoint Metrics

Metric Name	Category	KPI
Sharepoint Foundation Active Threads	SharePoint Server	True
Sharepoint Foundation Current Page Requests	SharePoint Server	False
Sharepoint Foundation Executing SQL Queries	SharePoint Server	False
Sharepoint Foundation Executing Time/Page Request	SharePoint Server	True
Sharepoint Foundation Incoming Page Requests Rate	SharePoint Server	False
Sharepoint Foundation Object Cache Hit Count	SharePoint Server	False
Sharepoint Foundation Reject Page Requests Rate	SharePoint Server	False
Sharepoint Foundation Responded Page Requests Rate	SharePoint Server	True
SQL query executing time	SharePoint Server	False
Network Received Data Rate	SharePoint Web Server	True
Network Sent Data Rate	SharePoint Web Server	True
Process Processor Time (%)	SharePoint Windows Service	False
Process Threads	SharePoint Windows Service	False

Oracle Weblogic Metrics

Metrics are collected for the Oracle Weblogic application service.

Table 4-37. Oracle Weblogic Metrics

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

Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

Table 4-38. Pivotal TC Server Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Pivotal TC Server	False
Buffer Pool<InstanceName> Memory Used	Pivotal TC Server	False
Buffer Pool<InstanceName> Total Capacity	Pivotal TC Server	False
Class Loading Loaded Class Count	Pivotal TC Server	False
Class Loading Total Loaded Class Count	Pivotal TC Server	False
Class Loading Unloaded Class Count	Pivotal TC Server	False

Table 4-38. Pivotal TC Server Metrics (continued)

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

Table 4-38. Pivotal TC Server Metrics (continued)

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

ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

Table 4-39. ActiveMQ Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Active MQ	False
Buffer Pool<InstanceName> Memory Used	Active MQ	False
Buffer Pool<InstanceName> Total Capacity	Active MQ	False
Class Loading Loaded Class Count	Active MQ	False
Class Loading Unloaded Class Count	Active MQ	False

Table 4-39. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
Class Loading Total Loaded Class Count	Active MQ	False
File Descriptor Usage Max File Descriptor Count	Active MQ	False
File Descriptor Usage Open File Descriptor Count	Active MQ	False
Garbage Collection<InstanceName> Total Collection Count	Active MQ	False
Garbage Collection<InstanceName> Total Collection Time	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Initial Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Used Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Initial Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Used Memory	Active MQ	False
Threading Thread Count	Active MQ	False
Uptime	Active MQ	False
UTILIZATION Process CpuLoad	Active MQ	False
UTILIZATION Memory Limit	ActiveMQ Broker	True

Table 4-39. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Memory Percent Usage (%)	ActiveMQ Broker	True
UTILIZATION Store Limit	ActiveMQ Broker	False
UTILIZATION Store Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION Temp Limit	ActiveMQ Broker	False
UTILIZATION Temp Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION Total Consumer Count	ActiveMQ Broker	True
UTILIZATION Total Dequeue Count	ActiveMQ Broker	True
UTILIZATION Total Enqueue Count	ActiveMQ Broker	True
UTILIZATION Total Message Count	ActiveMQ Broker	True
JVM Memory Heap Memory Usage Initial Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Used Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Initial Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Used Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Object Pending FinalizationCount	ActiveMQ JVM Memory Usage	False

Table 4-39. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Process CpuLoad	ActiveMQ OS	False
UTILIZATION System Cpu Load	ActiveMQ OS	False
UTILIZATION Consumer Count	ActiveMQ Topic	True
UTILIZATION Dequeue Count	ActiveMQ Topic	True
UTILIZATION Enqueue Count	ActiveMQ Topic	True
UTILIZATION Queue Size	ActiveMQ Topic	True
UTILIZATION Producer Count	ActiveMQ Topic	False

Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

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

Table 4-40. Apache HTTPD Metrics

Metric Name	Category	KPI
UTILIZATION Busy Workers	Apache HTTPD	True
UTILIZATION Bytes Per Req	Apache HTTPD	False
UTILIZATION Bytes Per Sec	Apache HTTPD	False
UTILIZATION CPU Load	Apache HTTPD	True
UTILIZATION CPU User	Apache HTTPD	False
UTILIZATION Idle Workers	Apache HTTPD	True
UTILIZATION Request Per Sec	Apache HTTPD	True
UTILIZATION SCBoard Closing	Apache HTTPD	False
UTILIZATION SCBoard DNS Lookup	Apache HTTPD	False
UTILIZATION SCBoard Finishing	Apache HTTPD	False
UTILIZATION SCBoard Idle Cleanup	Apache HTTPD	False
UTILIZATION SCBoard Keep Alive	Apache HTTPD	False
UTILIZATION SCBoard Logging	Apache HTTPD	False
UTILIZATION SCBoard Open	Apache HTTPD	False

Table 4-40. Apache HTTPD Metrics (continued)

Metric Name	Category	KPI
UTILIZATION SCBoard Reading	Apache HTTPD	False
UTILIZATION SCBoard Sending	Apache HTTPD	False
UTILIZATION SCBoard Starting	Apache HTTPD	False
UTILIZATION SCBoard Waiting	Apache HTTPD	False
UTILIZATION Total Accesses	Apache HTTPD	False
UTILIZATION Total Bytes	Apache HTTPD	True
UTILIZATION Total Connections	Apache HTTPD	False
UTILIZATION Uptime	Apache HTTPD	True
UTILIZATION Asynchronous Closing Connections	Apache HTTPD	False
UTILIZATION Asynchronous Keep Alive Connections	Apache HTTPD	False
UTILIZATION Asynchronous Writing Connections	Apache HTTPD	False
UTILIZATION ServerUptimeSeconds	Apache HTTPD	False
UTILIZATION Load1	Apache HTTPD	False
UTILIZATION Load5	Apache HTTPD	False
UTILIZATION ParentServerConfigGeneration	Apache HTTPD	False
UTILIZATION ParentServerMPMGeneration	Apache HTTPD	False

Oracle Database Metrics

Metrics are collected for the Oracle database application service.

Oracle database cannot be activated on Linux platforms.

Table 4-41. Oracle Database Metrics

Metric Name	Category	KPI
Utilization Active Sessions	OracleDB	True
Utilization Buffer CacheHit Ratio	OracleDB	False
Utilization Cursor CacheHit Ratio	OracleDB	False
Utilization Database Wait Time	OracleDB	False
Utilization Disk Sort persec	OracleDB	False
Utilization Enqueue Timeouts Persec	OracleDB	False

Table 4-41. Oracle Database Metrics (continued)

Metric Name	Category	KPI
Utilization Global Cache Blocks Corrupted	OracleDB	False
Utilization Global Cache Blocks Lost	OracleDB	False
Utilization Library CacheHit Ratio	OracleDB	False
Utilization Logon persec	OracleDB	True
Utilization Memory Sorts Ratio	OracleDB	True
Utilization Rows persort	OracleDB	False
Utilization Service Response Time	OracleDB	False
Utilization Session Count	OracleDB	True
Utilization Session Limit	OracleDB	False
Utilization Shared Pool Free	OracleDB	False
Utilization Temp Space Used	OracleDB	False
Utilization Total Sorts persec	OracleDB	False
Utilization Physical Read Bytes Persc	OracleDB	False
Utilization Physical Read IO Requests Persc	OracleDB	False
Utilization Physical Read Total Bytes Persec	OracleDB	False
Utilization Physical Reads Persec	OracleDB	True
Utilization Physical Reads Per Txn	OracleDB	False
Utilization Physical Write Bytes Persc	OracleDB	False
Utilization Physical Write IO Requests Persc	OracleDB	False
Utilization Physical Write Total Bytes Persc	OracleDB	False
Utilization Physical Writes Persc	OracleDB	True
Utilization Physical Writes Per Txn	OracleDB	False
Utilization User Commits Percentage	OracleDB	False
Utilization User Commits Persc	OracleDB	False
Utilization User Rollbacks Percentage	OracleDB	False
Utilization User Rollbacks persec	OracleDB	True

Table 4-41. Oracle Database Metrics (continued)

Metric Name	Category	KPI
Utilization User Transaction Persec	OracleDB	False
Utilization Database Time Persc	OracleDB	False

Cassandra Database Metrics

Metrics are collected for the Cassandra database application service.

Table 4-42. Cassandra Database Metrics

Metric Name	Category	KPI
Cache<InstanceName> Capacity	Cassandra	False
Cache<InstanceName> Entries	Cassandra	True
Cache<InstanceName> HitRate	Cassandra	True
Cache<InstanceName> Requests	Cassandra	True
Cache<InstanceName> Size	Cassandra	False
ClientRequest<InstanceName> Failures	Cassandra	False
ClientRequest<InstanceName> Latency	Cassandra	False
ClientRequest<InstanceName> Timeouts	Cassandra	False
ClientRequest<InstanceName> Total Latency	Cassandra	False
ClientRequest<InstanceName> Unavailables	Cassandra	False
CommitLog Pending Tasks	Cassandra	False
CommitLog Total Commit Log Size	Cassandra	False
Compaction Bytes Compacted	Cassandra	False
Compaction Completed Tasks	Cassandra	False
Compaction Pending Tasks	Cassandra	False
Compaction Total Compactions Completed	Cassandra	False
Connected Native Clients	Cassandra	False
HeapMemoryUsage committed	Cassandra	False
HeapMemoryUsage init	Cassandra	False
HeapMemoryUsage max	Cassandra	False

Table 4-42. Cassandra Database Metrics (continued)

Metric Name	Category	KPI
HeapMemoryUsage used	Cassandra	False
NonHeapMemoryUsage committed	Cassandra	False
NonHeapMemoryUsage init	Cassandra	False
NonHeapMemoryUsage max	Cassandra	False
NonHeapMemoryUsage used	Cassandra	False
ObjectPendingFinalizationCount	Cassandra	False
Storage Exceptions Count	Cassandra	False
Storage Load Count	Cassandra	False
Table<InstanceName> Coordinator Read Latency	Cassandra	False
Table<InstanceName> Live Diskspace Used	Cassandra	False
Table<InstanceName> Read Latency	Cassandra	False
Table<InstanceName> Total Diskspace Used	Cassandra	False
Table<InstanceName> Total Read Latency	Cassandra	False
Table<InstanceName> Total Write Latency	Cassandra	False
Table<InstanceName> Write Latency	Cassandra	False
ThreadPools<InstanceName> Active Tasks	Cassandra	False
ThreadPools<InstanceName> Currently Blocked Tasks	Cassandra	False
ThreadPools<InstanceName> Pending Tasks	Cassandra	False

Hyper-V Metrics

Metrics are collected for the Hyper-V application service.

Table 4-43. Hyper-V Metrics

Metric Name	Category	KPI
VM:Hyper-V Virtual Machine Health Summary Health Critical	HyperV	False
VM<instanceName> Physical Memory	HyperV	False
VM<instanceName> Hv VP 0 Total Run Time	HyperV	False

Table 4-43. Hyper-V Metrics (continued)

Metric Name	Category	KPI
VM<instanceName> Bytes Received	HyperV	False
VM<instanceName> Bytes Sent	HyperV	False
VM<instanceName> Error Count	HyperV	False
VM<instanceName> Latency	HyperV	False
VM<instanceName> Queue Length	HyperV	False
VM<instanceName> Throughput	HyperV	False
CPU<instanceName> Idle Time	HyperV	True
CPU<instanceName> Processor Time	HyperV	True
CPU<instanceName> User Time	HyperV	True
Disk<instanceName> Avg Disk Queue Length	HyperV	False
Disk<instanceName> Idle Time	HyperV	False
Disk<instanceName> Read Time	HyperV	True
Disk<instanceName> Write Time	HyperV	True
Process<instanceName> Private Bytes	HyperV	False
Process<instanceName> Processor Time	HyperV	False
Process<instanceName> Thread Count	HyperV	False
Process<instanceName> User Time	HyperV	False
System Processes	HyperV	False
System Processor Queue Length	HyperV	False
System System UpTime	HyperV	False
Memory Available Bytes	HyperV	False
Memory Cache Bytes	HyperV	False
Memory Cache Faults	HyperV	False
Memory Pages	HyperV	False
Network<instanceName> Packets Outbound Error	HyperV	False
Network<instanceName> Packets Received Error	HyperV	False

MongoDB Metrics

Metrics are collected for the MongoDB application service.

Table 4-44. MongoDB Metrics

Metric Name	Category	KPI
UTILIZATION Active Reads	MongoDB	True
UTILIZATION Active Writes	MongoDB	True
UTILIZATION Connections Available	MongoDB	False
UTILIZATION Connections Total Created	MongoDB	False
UTILIZATION Current Connections	MongoDB	True
UTILIZATION Cursor Timed Out	MongoDB	True
UTILIZATION Deletes Per Sec	MongoDB	False
UTILIZATION Document Inserted	MongoDB	False
UTILIZATION Document Deleted	MongoDB	False
UTILIZATION Flushes Per Sec	MongoDB	False
UTILIZATION Inserts Per Sec	MongoDB	False
UTILIZATION Net Input Bytes	MongoDB	False
UTILIZATION Open Connections	MongoDB	True
UTILIZATION Page Faults Per Second	MongoDB	False
UTILIZATION Net Output Bytes	MongoDB	False
UTILIZATION Queries Per Sec	MongoDB	False
UTILIZATION Queued Reads	MongoDB	True
UTILIZATION Queued Writes	MongoDB	True
UTILIZATION Total Available	MongoDB	False
UTILIZATION Total Deletes Per Sec	MongoDB	False
UTILIZATION Total Passes Per Sec	MongoDB	False
UTILIZATION Total Refreshing	MongoDB	False
UTILIZATION Updates Per Sec	MongoDB	False
UTILIZATION Volume Size MB	MongoDB	False
UTILIZATION Collection Stats	MongoDB DataBases	False
UTILIZATION Data Index Stats	MongoDB DataBases	True

Table 4-44. MongoDB Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Data Indexes	MongoDB DataBases	False
UTILIZATION Data Size Stats	MongoDB DataBases	True
UTILIZATION Average Object Size stats	MongoDB DataBases	False
UTILIZATION Num Extents Stats	MongoDB DataBases	False

Riak Metrics

Metrics are collected for the Riak application service.

Table 4-45. Riak Metrics

Metric Name	Category	KPI
UTILIZATION CPU Average	Riak KV	False
UTILIZATION Memory Processes	Riak KV	False
UTILIZATION Memory Total	Riak KV	False
UTILIZATION Node GETs	Riak KV	True
UTILIZATION Node GETs Total	Riak KV	False
UTILIZATION Node PUTs	Riak KV	True
UTILIZATION Node PUTs Total	Riak KV	False
UTILIZATION PBC Active	Riak KV	True
UTILIZATION PBC Connects	Riak KV	True
UTILIZATION Read Repairs	Riak KV	True
UTILIZATION vNODE Index Reads	Riak KV	True
UTILIZATION vNODE Index Writes	Riak KV	True

NTPD Metrics

Metrics are collected for the NTPD application service.

Table 4-46. NTPD Metrics

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

Table 4-46. NTPD Metrics (continued)

Metric Name	Category	KPI
ntpd reach	Network Time Protocol	True
ntpd when	Network Time Protocol	False

WebSphere Metrics

Metrics are collected for the WebSphere application service.

Table 4-47. WebSphere Metrics

Metric Name	Category	KPI
Thread Pool Active Count Current	Thread Pool	False
Thread Pool Active Count High	Thread Pool	False
Thread Pool Active Count Low	Thread Pool	False
Thread Pool Active Count Lower	Thread Pool	False
Thread Pool Active Count Upper	Thread Pool	False
JDBC Close Count	JDBC	False
JDBC Create Count	JDBC	False
JDBC JDBC Pool Size Average	JDBC	False
JDBC JDBC Pool Size Current	JDBC	False
JDBC JDBC Pool Size Lower	JDBC	False
JDBC JDBC Pool Size Upper	JDBC	False
Garbage Collection<InstanceName> Total Collection Count	WebSphere	False
Garbage Collection<InstanceName> Total Collection Time	WebSphere	False
JVM Memory Heap Memory Usage Committed Memory	WebSphere	False
JVM Memory Heap Memory Usage Initial Memory	WebSphere	False

Table 4-47. WebSphere Metrics (continued)

Metric Name	Category	KPI
JVM Memory Heap Memory Usage Maximum Memory	WebSphere	False
JVM Memory Heap Memory Usage Used Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Committed Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Initial Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Maximum Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Used Memory	WebSphere	False
JVM Memory Number of Object Pending Finalization Count	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Committed Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Initial Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Used Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Committed Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Initial Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Maximum Memory	WebSphere	False

Table 4-47. WebSphere Metrics (continued)

Metric Name	Category	KPI
JVM Memory Pool<InstanceName> Usage Used Memory	WebSphere	False
Process Cpu Load	WebSphere	False
System Cpu Load	WebSphere	False
System Load Average	WebSphere	False

Java Application Metrics

Metrics are collected for the Java application service.

Table 4-48. Java Application Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Java Application	False
Buffer Pool<InstanceName> Memory Used	Java Application	False
Buffer Pool<InstanceName> Total Capacity	Java Application	False
Class Loading Loaded Class Count	Java Application	True
Class Loading Total Loaded Class Count	Java Application	False
Class Loading Unloaded Class Count	Java Application	False
Garbage Collection<InstanceName> Total Collection Count	Java Application	False
Garbage Collection<InstanceName> Total Collection Time	Java Application	False
JVM Memory Heap Memory Usage Committed Memory	Java Application	False
JVM Memory Heap Memory Usage Initial Memory	Java Application	False
JVM Memory Heap Memory Usage Maximum Memory	Java Application	False
JVM Memory Heap Memory Usage Used Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Committed Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Initial Memory	Java Application	False

Table 4-48. Java Application Metrics (continued)

Metric Name	Category	KPI
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Used Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Committed Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Initial Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Maximum Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Used Memory	Java Application	False
JVM Memory Non Heap Memory Usage Committed Memory	Java Application	False
JVM Memory Non Heap Memory Usage Initial Memory	Java Application	False
JVM Memory Non Heap Memory Usage Maximum Memory	Java Application	False
JVM Memory Non Heap Memory Usage Used Memory	Java Application	False
JVM Memory Object Pending Finalization Count	Java Application	False
Uptime	Java Application	True
Threading Thread Count	Java Application	True
Process CPU Usage %	Java Application	False
System CPU Usage %	Java Application	False
System Load Average %	Java Application	False

Remote Check Metrics

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

HTTP Metrics

vRealize Operations Manager discovers metrics for HTTP remote checks.

HTTP Metrics

Table 4-49. HTTP Metrics

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

ICMP Metrics

vRealize Operations Manager discovers metrics for the ICMP object type.

Table 4-50. ICMP Metrics

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

TCP Metrics

vRealize Operations Manager discovers metrics for the TCP object type.

Table 4-51. TCP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

UDP Metrics

vRealize Operations Manager discovers metrics for the UDP object type.

Table 4-52. UDP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

Troubleshooting

Troubleshooting the Configuration of vRealize Application Remote Collector

vRealize Application Remote Collector Configuration Fails

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

Problem

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

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

Solution

- ◆ Enable the relevant ports.
- ◆ Ensure that vRealize Operations Manager and vRealize Application Remote Collector have the NTP synced.

Error in Password

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

Problem

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

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

Solution

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

Troubleshooting Agent Installation

Agent Install Failure Because of the vCenter Server User Permissions

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

Problem

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

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

Solution

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

Agent Install Failure Because NTP is Not in Sync

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

Problem

- Agent installation fails
- Adapter configuration fails

Solution

- ◆ Ensure that you configure network time protocol settings, or
- ◆ Run the following command to update the time immediately from an NTP server: `ntpdate time.vmware.com`

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

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

Agent Install Fails on a Linux End Point

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

Problem

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

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

Solution

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

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

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

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

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

Agent Install on Windows Fails When UAC is Disabled

Problem

Install of the agent fails even when UAC is disabled.

Solution

- ◆ To disable UAC (previously known as LUA) on Windows, complete the following steps:
 - In the registry path `HKLM:\Software\Microsoft\Windows\CurrentVersion\Policies\System`, set the value for the key `EnableLUA` to `0`.
 - You must reboot the machine for the changes to take effect.

Agent Install Fails on Windows with a Permission Denied Error

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

Problem

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

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

Solution

- ◆ Disable the antivirus and then proceed with bootstrapping.

Troubleshooting Plugin Related Failures

Unable to Activate a Plugin

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

Problem

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

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

Solution

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

Plugin Status Is Displayed as Unknown

The status of a few plugins is Unknown after vRealize Application Remote Collector and vRealize Operations Manager are upgraded from 7.5 to 8.3 and 8.1 to 8.3.

Problem

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

Solution

- ◆ Reactivate the plugin.

Troubleshooting Metric Collection

Troubleshoot Agent Installation and Metric Collection Issues

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

Problem

You might notice the following issues in vRealize Operations Manager:

- You cannot add vRealize Application Remote Collector to vRealize Operations Manager
- You cannot install an agent in the Windows and Linux target VMs.

Cause

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

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

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

Solution

- 1 Check the vRealize Operations Manager support bundle in the following path: `COLLECTOR/adapters/APPOSUCPAdapter/` for errors.
- 2 Check the vRealize Application Remote Collector support bundle, *ucpapi.log*, for errors.
- 3 Ensure time synchronization between vRealize Application Remote Collector, vRealize Operations Manager and the Windows and Linux target VMs.
- 4 To start and restart the agent, see [Additional Operations from the Manage Agents Tab](#).

Troubleshooting Upgrade

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

Problem

vRealize Application Remote Collector UI Problems

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

Manage vRealize Application Remote Collector UI Problems

- You can see an option to update the endpoint agent but you are unable to perform the update.
- Services supported in the latest versions of vRealize Application Remote Collector cannot be discovered.

Cause

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

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

Solution

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

Note For more troubleshooting steps, see [Troubleshooting Agent Installation](#).

Troubleshooting Content Upgrade

Problem

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

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

Cause

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

Solution

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

Troubleshooting Using Support Bundles

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

For vRealize Application Remote Collector

- 1 Access the VAMI page by entering `https://<vRealize Application Remote Collector hostname>:5480`
- 2 Log in with root credentials.
- 3 Click the **Support Bundle** tab. Click the **Generate Logs for VA** button.
vRealize Application Remote Collector creates the support bundles which you can download.

For End Point VMs

- 1 Log in to the end point.
- 2 Run the following commands based on the end point VM's operating system type:

For Linux End Point VMs

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

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

For Windows End Point VMs

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

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

Service Discovery

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

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

Licensing

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

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

- Configure Service Discovery. For more information, see [Configure Service Discovery](#).
- Manage Services. For more information, see [Manage Services](#).
- Monitor services using dashboards. For more information, see [Service Discovery Dashboards](#).
- View the services discovered. For more information, see [Discovered Services](#).

Supported Platforms and Products for Service Discovery

Service discovery supports specific platforms and product versions.

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

Supported Product Versions for Credential-Based Service Discovery

- For ESXi, vCenter Server, and VMware Cloud on AWS versions, see the [VMware Product Interoperability Matrix](#).
- VMware Tools: For details, see [KB 75122](#).

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

For information, see [KB 78216](#).

Operating System Versions

Operating Systems	Version
Windows	Windows 7, Windows Server 2008/R2, and above.
Linux	Photon, RHEL, CentOS, SUSE Linux Enterprise Server, OEL, and Ubuntu (all Linux operating systems must be based on kernel version 2.6.25 or above).

Supported Services

Service discovery supports several services that are supported in vRealize Operations Manager. The supported services are listed here.

Supported Services:

- Active Directory
- Apache HTTP
- Apache Tomcat
- DB2
- Exchange Client Access Server

- Exchange Edge Transport Server
- Exchange Hub Transport Server
- Exchange Mailbox Server
- Exchange Server
- Exchange Unified Messaging Server
- GemFire
- IIS
- JBoss
- MS SQL DB
- MySQL DB
- Nginx
- Oracle DB
- RabbitMQ
- SharePoint
- SharePoint Application Server
- SharePoint Server
- SharePoint Web Server
- SRM vCenter Replication Management Server
- SRM vCenter Replication Server
- Sybase DB
- Pivotal tc Server
- vCenter Site Recovery Manager Server
- vCloud Director
- VMware vCenter
- VMware vCenter (Appliance)
- VMware View Server
- vRealize Operations Analytics
- vRealize Operations Collector
- vRealize Operations GemFire
- vRealize Operations Postgres Data
- vRealize Operations Postgres Repl
- vRealize Operations UI

- WebLogic
- WebSphere

Configure Service Discovery

To discover services and their relationships and to access basic monitoring, you can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

Prerequisites

- You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to discover services. The configured vCenter Server user must have the following privileges:
 - Guest operation alias modification
 - Guest operation alias query
 - Guest operation modifications
 - Guest operation program execution
 - Guest operation queries
 - Manage service configurations
 - Modify service configuration
 - Query service configuration
 - Read service configuration
- The ESXi instance that hosts the VMs where services should be discovered, must have HTTPS access to port 443 from the collector node on which the service discovery adapter instance is configured.
- Verify that the following types of commands and utilities are used:

Type	Commands and Utilities
UNIX Operating Systems	
Service Discovery	<code>ps</code> , <code>netstat</code> , and <code>top</code>
Performance Metrics Collection	<code>:</code> <code>awk</code> , <code>csh</code> , <code>ps</code> , <code>pgrep</code> , and <code>procfs</code> (file system)
Windows Operating Systems	
Service Discovery	<code>wmic</code> and <code>netstat</code>
Performance Metrics Collection	<code>wimic</code> , <code>typeperf</code> , and <code>tasklist</code>

■ User Access Restrictions

- For Linux operating systems, ensure that the user is a root or member of the *sudo* users group.

Note For non-root users, the `NOPASSWD` option must be enabled in `/etc/sudoers` file to avoid the metrics collector scripts from waiting for the interactive password input.

Steps to enable the `NOPASSWD` option for a particular sudo user:

- 1 Login to the specific VM as a root user.
 - 2 Run the `sudo visudo` command that opens an editor.
 - 3 In the command section, add `username ALL=(ALL) NOPASSWD:ALL`. `username` must be replaced with an existing user name for which this option is enabled.
 - 4 Save the file and close it. It is automatically reloaded.
-

- To discover services on Windows, the local administrator account must be configured.

Note Services will not be discovered for administrator group members that are different from the administrator account itself if the policy setting `User Account Control: Run all administrators in Admin Approval Mode` is turned on. As a workaround, you can turn off this policy setting to discover services. However, if you turn the policy setting off, the security of the operating system is reduced.

- To discover services on Windows Active Directory, the domain administrator account must be configured.
- The system clock must be synchronized between the vRealize Operations Manager nodes, the vCenter Server, and the VM if service discovery is working in credential-based mode and guest alias mapping is used for authentication.
- The configured user must have read and write privileges to the temp directory. For Windows systems, the path can be taken from the environment variable `TEMP`. For Linux systems, it is `/tmp` and/or `/var/tmp`.
- For more information about supported platforms and versions, see [Supported Platforms and Products for Service Discovery](#).

Note If more than one vRealize Operations Manager instance is monitoring the same vCenter Server and service discovery is enabled for those vRealize Operations Manager instances, then service discovery might be unstable, which is a known VMware Tools problem. As a result, guest operations might fail to execute.

Procedure

- 1 In the menu, select **Home** and then select **Manage Applications > Discover Services** from the left panel.
- 2 From the **Discover Services** page, click the **Configure Service Discovery** option.

- 3 From the **Cloud Accounts** page, click the vCenter Server instance from the list and then select the **Service Discovery** tab.
- 4 To enable service discovery in this vCenter Server, enable the **Service Discovery** option.
- 5 You can choose to add credentials by selecting the **Use alternate credentials** check box.
 - a Click the plus sign and enter the details in the **Manage Credentials** dialog box, which include a credential name and a vCenter user name and password. In addition, enter the user name and password for Windows, Linux, and SRM and click **OK**.
- 6 Alternatively, if you are using the default user name and password, enter a default user name and password for Windows, Linux, and SRM.
- 7 Enter a password for the guest user mapping.
- 8 You can also enable grouping of the application and the creation of a business application.
- 9 Click **Save**.

Note If you specify a non-root user for Linux, services are not discovered unless you enable the option Use Sudo (Linux Non-root user) while editing the associated Service Discovery adapter instance after you create the vCenter Cloud Account. This option is disabled by default, which means the root user is expected by default when you configure the vCenter Cloud Account.

- 10 Edit the cloud account created for service discovery.
- 11 In the **Advanced Settings** section, to configure credential-less service discovery, select **Enabled** from the **Credential-less service discovery status** field.

What to do next

You can manage services supported by vRealize Operations Manager on specific VMs.

Manage Services

You can manage services supported by vRealize Operations Manager on the specific VMs.

Where You Manage Services

In the menu, select **Administration** and then select **Inventory** from the left panel. Select the **Manage Services** tab from the right pane. You can also navigate to the **Manage Services** tab by selecting **Home**, and then select **Manage Applications > Discover Services** from the left pane. Select the **Manage Services** option from the **Discover Services** page.

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

Table 4-53. Datagrid Options

Options	Description
VM Name	Name of the VM.
Operating System	Operating system installed on the VM.

Table 4-53. Datagrid Options (continued)

Options	Description
Services Discovered	Displays the names of discovered services or <i>None</i> , if services are not discovered on the VM.
Service Monitoring	Displays the current value of the VM's service monitoring setting. If set, services are discovered and service performance metrics are calculated every 5 minutes. Otherwise, service discovery is performed every 24 hours.
Authentications Status	<p>VM authentication status for service discovery. The possible values are:</p> <ul style="list-style-type: none"> ■ Unknown ■ Failed ■ Guest Alias ■ Common Credentials ■ Credential-less
Power State	<p>Power status of the VMs. The possible values are:</p> <ul style="list-style-type: none"> ■ Powered On ■ Powered Off ■ Suspended ■ Unknown
Collection State	Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the collection state icon. To manage an adapter instance to start and stop collection of data, in the menu, click Administration , and then in the left pane click Inventory .
Collection Status	<p>Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the collection status icon. To manage an adapter instance to start and stop collection of data, in the menu, click Administration, and then in the left pane click Inventory.</p> <p>You can view a message for VMs with a failed authentication status in a tool tip when you point to the collection status icon.</p>
vCenter Name	Name of the vCenter Adapter instance to which that VM resource belongs.

Table 4-54. Toolbar Options

Options	Description
Actions	Displays a list of actions. For more information, see List of vRealize Operations Manager Actions .
Open in external application	If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application.

Table 4-54. Toolbar Options (continued)

Options	Description
Provide Password	Select VMs from the list and click Provide Password to provide a user name and password for the selected VMs to discover the services.
Enable Service Monitoring	Select VMs from the list and click Enable Service Monitoring to enable frequent service discovery and service performance metrics calculation (every 5 minutes). Note Selecting too many VMs will potentially result in vCenter Server degradation which is a known issue.
Disable Service Monitoring	Select VMs from the list and click Disable Service Monitoring to disable frequent service discovery and service performance metrics calculation. Service discovery defaults to the 24-hour cycle.
Clear Selections	Clears all VM object selections.
Select All	Selects all VM objects.
Show Detail	Navigates to the Summary tab for the selected VM.
Page Size	The number of objects to list per page.
All Filters	You can search through the list of VMs according to the following criteria: VM Name, Operating System, Power State, Status, and Service.

Discovered Services

You can view discovered services, the number of VMs on which each discovered service is running, and you can configure service discovery.

Where You View the Discovered Services

From the menu, select **Home**, and then from the left pane select **Discover Services**.

Discovered Services

You see a list of services that are discovered and the number of VMs that have the services running. You see this section after you have configured Service Discovery and the services are discovered.

Known Services

You see a list of all the services supported and those that can be discovered.

Allowed Services

You can configure a service by clicking **Configure Allowed List**, and adding a process name, port, and display name in the **Allow Service** dialog box.

The process name must exactly match the name that you see in the guest OS when running commands `ps` in Linux and `wmic` in Windows. Specify a single port for each service.

Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

Table 4-55. Virtual Machine Metrics

Metric Name	Description
Guest OS Services Total Number of Services	Number of out-of-the-box and user-defined services discovered in the VM.
Guest OS Services Number of User Defined Services	Number of user-defined services discovered in the VM.
Guest OS Services Number of OOTB Services	Number of out-of-the-box services discovered in the VM.
Guest OS Services Number of Outgoing Connections	Number of outgoing connection counts from the discovered services.
Guest OS Services Number of Incoming Connections	Number of incoming connection counts to the discovered services.

Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

Table 4-56. Service Summary Metrics

Metric Name	Description
Summary Incoming Connections Count	Number of incoming connections.
Summary Outgoing Connections Count	Number of outgoing connections.
Summary Connections Count	Number of incoming and outgoing connections.
Summary Pid	Process ID.

Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

Table 4-57. Service Performance Metrics

Metric Name	Description
Performance metrics group CPU	CPU usage in percentage.
Performance metrics group Memory	Memory usage in KB.

Table 4-57. Service Performance Metrics (continued)

Metric Name	Description
Performance metrics group IO Read Throughput	IO read throughput in KBps.
Performance metrics group IO Write Throughput	IO write throughput in KBps.

Service Type Metrics

Service discovery discovers metrics for service type objects.

Table 4-58. Service Type Metrics

Metric Name	Description
Number of instances	Number of instances of this service type.

Log Insight

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

Log Insight Page

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

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

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

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

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

Logs Tab

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

How the Logs Tab Works

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

Where You Find the Logs Tab

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

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

Configuring vRealize Log Insight with vRealize Operations Manager

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

Configuring the vRealize Log Insight Adapter in vRealize Operations Manager

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

You can integrate only one vRealize Log Insight instance.

Prerequisites

- Verify that vRealize Log Insight and vRealize Operations Manager are installed.
- Verify that you know the IP address or FQDN of the vRealize Log Insight instance you have installed.

Procedure

- 1 In the menu, select **Administration**, and then from the left pane, select **Management > Integrations**.
- 2 From the **Integrations** page, click VMware vRealize Log Insight.
- 3 In the VMware vRealize Log Insight page complete the following steps:
 - Enter the IP address or FQDN in the **Log Insight server** text box of the vRealize Log Insight you have installed and want to integrate with.
 - Select the collector group from the **Collectors/Groups** drop-down menu.
 - Click **Test Connection** to verify that the connection is successful.
 - Click **Save**.

- 4 From the vRealize Operations Manager Home page, click **Troubleshoot > Using Logs** from the left pane. If you see a statement at the bottom of the page, click the link and accept the certificate exception in vRealize Log Insight or contact your IT support for more information.
- 5 From the vRealize Operations Manager Home page, click **Troubleshoot > Using Logs** from the left pane and enter the user name and password of the vRealize Log Insight instance you have installed.

Configuring vRealize Operations Manager in vRealize Log Insight

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

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

Prerequisites

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

Procedure

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

Note If you are using a load balancer, use its IP address or FQDN as a hostname value.

- 3 In the **Username** and **Password** text boxes, enter the user name and password of the vRealize Operations Manager instance you want to integrate with.
- 4 Select the relevant check boxes according to your preference:
 - To send alerts to vRealize Operations Manager, select **Enable alerts integration**.
 - To let vRealize Operations Manager open Log Insight and query for object logs, select **Enable launch in context**.
 - To calculate and send metrics to vRealize Operations Manager, select **Enable metric calculation**.
- 5 Click **Test Connection** to verify that the connection is successful and accept the certificate if it is untrusted.
- 6 Click **Save**.

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

Log Forwarding

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

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

Where You Find the Log Forwarding Page

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

Table 4-59. Log Forwarding Page Options

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

Table 4-59. Log Forwarding Page Options (continued)

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

Modifying Existing Log Types

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

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

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

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

```
vmw_vr_ops_appname
vmw_vr_ops_clustername
vmw_vr_ops_clusterrole
vmw_vr_ops_hostname
vmw_vr_ops_nodename
```

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

Business Management

SDDC costing is out-of-the box with vRealize Operations Manager . There is no integration required with vRealize Business for Cloud.

Cost Settings for Financial Accounting Model

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

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

Deprecated Metrics

The MTD CPU Cost, MTD Memory Cost, MTD Storage Cost, VM Direct Cost have been deprecated since vRealize Operations Manager . The corresponding metrics to refer instead of these metrics would be the daily cost metrics.

Configuring Depreciation Preferences

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

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

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

Example: Example for Straight Line Depreciation Method

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

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

Year	Original Cost	Depreciation Rate	Accumulated Depreciation	Straight Line Depreciation Cost
Year 1	10000	0.4	0	$\text{Maximum}([(10000-0) * 0.4], [(10000-0) / 5])$ $= \text{Maximum}(4000, 2000) = 4000$ <p>which is 333.33 per month.</p>
Year 2	10000	0.4	4000	$\text{Maximum}([(10000-4000) * 0.4], [(10000-4000) / 4])$ $= \text{Maximum}(2400, 1500) = 2400$ <p>which is 200 per month.</p>
Year 3	10000	0.4	6400	$\text{Maximum}([(10000-6400) * 0.4], [(10000-6400) / 3])$ $= \text{Maximum}(1440, 1200) = 1440$ <p>which is 120 per month.</p>
Year 4	10000	0.4	7840	$\text{Maximum}([(10000-7840) * 0.4], [(10000-7840) / 2])$ $= \text{Maximum}(864, 1080) = 1080$ <p>which is 90 per month.</p>
Year 5	10000	0.4	8920	$\text{Maximum}([(10000-8920) * 0.4], [(10000-8920) / 1])$ $= \text{Maximum}(432, 1080) = 1080$ <p>which is 90 per month.</p>

Overview of Cost Drivers

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

You can now set a total cost for the License, Labor, Network, Maintenance, and facilities cost drivers in vRealize Operations Manager :

Note The total cost set by you is distributed across resources in the data center. For example, if you set the total cost for the RHEL license, the cost is divided across all the hosts and VMs which use the RHEL license.

According to the industry standard, vRealize Operations Manager maintains a reference cost for these cost drivers. This reference cost helps you for calculating the cost of your setup, but might not be accurate. For example, you might have received some special discounts during a bulk purchase or you might have an ELA with VMware that might not match the socket-based pricing available in the reference database. To get accurate values, you can modify the reference cost of cost drivers in vRealize Operations Manager , which overrides the values in the reference database. Based on your inputs, vRealize Operations Manager recalculates the total amount for

the private cloud expenses. After you add a private cloud into vRealize Operations Manager , vRealize Operations Manager automatically discovers one or more vCenter Servers that are part of your Private Cloud. In addition, it also retrieves the inventory details from each vCenter Server. The details include:

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

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

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

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

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

Table 4-60. Expense Types

Cost Drivers	Description
Server Hardware : Traditional	<p>The Server Hardware cost driver tracks all the expenses for purchasing of hardware servers that are part of vCenter Servers. You see the server cost based on CPU age and server cost details.</p> <p>Note You can now select an individual server from the server group and specify the unique cost for each individual server.</p>
Server Hardware : Hyper-Converged	<p>The Server Hardware : Hyper-Converged cost driver, tracks the expenses associated with hyper converged infrastructure components. The Server Hardware : Hyper-Converged cost driver includes expenses for the Hyper Converged servers like vSAN enabled servers and vXRail. The expense provided is for both compute and storage.</p> <p>Note The customizations that were performed for vSAN server costing under Server Hardware : Traditional in the earlier versions will not be carried forward to 7.5 as the vSAN enabled servers will fall under Server Hardware : Hyper-Converged servers now.</p>
Storage	<p>You can calculate the storage cost at the level of a datastore based on the tag category information collected from vCenter Server. You see the storage total distribution based on category and the uncategorized cost details.</p> <p>Note The vSAN datastores are not displayed as part of this cost driver page.</p>

Table 4-60. Expense Types (continued)

Cost Drivers	Description
License	<p>You see the licenses cost distribution for the operating systems cost and VMware license of your cloud environment.</p> <p>Note For Non-ESX physical servers, VMware license is not applicable.</p>
Maintenance	<p>You see the maintenance cost distribution for the server hardware and operating system maintenance. You can track your total expense with hardware and operating system vendors.</p>
Labor	<p>You see the labor cost distribution for the servers, virtual infrastructure, and operating systems. You can view the total administrative cost for managing physical servers, operating systems and virtual machines. You can track all expenses spent on human resources to manage the datacenters.</p> <p>Note</p> <ul style="list-style-type: none"> ■ Labor cost includes expenses on backup appliance virtual machine (VDP virtual appliance). ■ For physical servers, operating system labor cost and servers labor costs are applicable, virtual infrastructure cost is not considered.
Network	<p>You see the networks costs by NIC type. You can track a network expense based on different types of NICs attached to the ESX server. You can view the total cost of physical network infrastructure that includes the internet bandwidth, and is estimated by count and type of network ports on the ESXi Servers.</p> <p>Note For physical servers, the network details are not captured. So, the network cost is considered as zero.</p>
Facilities	<p>You see the cost distribution for the facilities such as real estate costs, such as rent or cost of data center buildings, power, cooling, racks, and associated facility management labor cost. You can point to the chart to see the cost details for each facility type.</p>
Additional Cost	<p>You can see the additional expenses such as backup and restore, high availability, management, licensing, VMware software licensing.</p>
Application Cost	<p>You can see the cost of different application services you are running in your environment compared to your overall expenses. Some examples of application cost are, cost of running SQL server cluster and cost of running Antivirus on VMs.</p>

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

Cloud Providers Overview

By default, you can see that Amazon Web Services (AWS), Google Cloud, IBM Cloud, and Microsoft Azure are included in vRealize Operations Manager . You can also add your own cloud provider by using a standard vRealize Operations Manager template.

You can configure the new cloud provider as per the standard vRealize Operations Manager template and perform a migration scenario. The vRealize Operations Manager template contains data points for vCPU, CPU, RAM, OS, region, plan term, location, and built-in instance storage, you must provide these values when you add cloud providers. The result of the migration scenario helps you assess the cost savings achieved using your cloud provider against the default cloud providers.

You can edit the rate card for new cloud providers and default cloud providers. However, you cannot delete the default cloud providers.

Add Cloud Provider

You can use the Add Cloud Provider workspace to add or edit a cloud provider. You can edit the cloud provider rate card for default cloud providers and the new cloud provider.

Procedure

- 1 On the menu, click **Administration** and in the left pane click **Configuration > Cost Settings > Cloud Providers**.

You can also reach the Cloud Providers page from the Home Screen. In the Home screen, navigate to **Optimize Capacity > What-If Analysis > Plan Migration > Add Cloud Providers**. For more information, see **What-If-Analysis - Migration Planning** section in vRealize Operations Manager help.

- 2 Click the **Add Cloud** icon.
- 3 Enter the **Cloud Provider Name**.
- 4 Select the cloud provider logo and click **Upload Logo**.
- 5 Click **Next**.
- 6 Click **Download Template** and specify the required values.

Note When you edit a cloud provider the Download Template link is replaced with Download Existing Rate Card. You can update the existing rate card and upload the same.

- 7 Select the updated template and click **Upload Rate Card**.
- 8 Click **Validate**.

Note vRealize Operations Manager validates the rate card and reports success or failure. If errors are reported, you can correct the errors and proceed further.

- 9 Click **Finish**.

Results

The new cloud provider is now part of the vRealize Operations Manager cloud provider list.

Billing Framework for Unmanaged Objects

You can remove objects which should not be monitored by vRealize Operations Manager using the billing framework. The billing framework ensures that the license fee is not applicable to the unmanaged objects which are moved to the maintenance state.

How To Manage Unmanaged Objects

To manage the unmanaged objects, you have to perform the following actions in vRealize Operations Manager :

- Remove objects that should not be monitored.
- Move the unmonitored objects to maintenance state.

- Stop the data collection for the objects in maintenance mode.
- Power off the virtual machines that are in maintenance mode.

Billing Support for Unmanaged Objects

When you remove specific objects from monitoring, vRealize Operations Manager moves these objects to maintenance mode and stops billing for the objects. The billing framework ensures that the costs related to licensing are not calculated for the following scenarios:

- vSphere and Public cloud virtual machines are in maintenance mode.
- vSphere and Public cloud virtual machines are in powered off state.
- vSphere and Public cloud virtual machines have stopped data collection.

The licensing fee is not charged for the objects in maintenance mode, you can verify the same in the next hourly billing cycle. You can navigate to **Administrator > Inventory** list, to view the list of objects that are in maintenance mode.

Billing Enhancements for Horizon Management Pack and Virtual Hosts

The cost calculation of vRealize Operations Manager has been enhanced to include the end point objects of Horizon Management Pack and virtual hosts. Earlier, the cost calculation was based on the metrics collected for each end point object.

The cost calculation for the end point objects is now based on the following criteria:

- Each Virtual Desktop Infrastructure Virtual Machine (VDI VM) is counted as 0.25 Operating System Instance (OSI)
- Each Remote Desktop Service Host (RDS Host) is counted as 0.25 Operating System Instance
- One Operating System Instance for each Connection Server
- Virtual Hosts (ESXi hosted on a VM) is not counted against license usage
- VMs hosting the virtual hosts are counted against license usage

There are no VDI VM objects discovered by Horizon MP. Instead, Horizon MP objects have relationships with vCenter MP Virtual Machines. VDI VMs are identified by their parent VDI Pool objects. vRealize Operations Manager for Cloud, reports the number of VDI VMs in the bill. The number of VDI VMs appear under Virtual Machine node of the vCenter MP.

How to Identify the Virtual Host

You can identify the virtual hosts by the following property.

- Hardware |Vendor = "VMware, Inc"

Editing Cost Drivers

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

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

Editing Server Hardware : Traditional

You can view, add, edit, or delete the cost of each server group, based on their configuration and the purchase date of a batch server running in your cloud environment. You can also specify the server cost for individual servers in a server group. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Server Hardware : Traditional**.

Note You can customize the default value of cost per server and specify exclusive values for other servers in the list.

For example, if you have a system that has eight servers you can modify the default reference value from \$1000 to \$800 for eight servers. You can also select two servers from the list and customize their value as \$600. So, any new server that is added to the system will have the default value as \$800.

- 3 Select the required edit mode for changing the server hardware cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
 - **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

- 4 Click any server from the list of **Server Group Description**.

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

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

5 After selecting a server group, you can manually enter the required fields.

- a Enter the Purchase Type and Cost Per Server.

Note You can use the **+ ADD COST PER SERVER** option to create multiple server batches and set the cost for a specific server in a server group.

- b Click **Save**.

Editing Server Hardware: Hyper-Converged

You can view, add, edit, or delete the cost of Hyper converged Infrastructure (HCI) component in your server group. You can specify the cost per server and compute percentage exclusively for the HCI servers. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

Procedure

1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.

2 In the Cost Drivers tab, click **Server Hardware : Hyper-Converged**.

3 Select the required edit mode for changing the server hardware cost.

- **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 Click any server from the list of **Server Group Description**.

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

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

Note You can edit the Compute Pct column to adjust the storage rate of the vSAN datastores. You can use the same percentage to determine the cost.

5 After selecting a server group, you can manually enter the required fields.

- a Enter Purchase Type, Cost Per Server, and Compute Percentage.

Note You can use the **+ ADD COST PER SERVER** option to create multiple server batches and to customize the cost per server.

- b Click **Save**.

Edit Monthly Cost of Storage

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

Prerequisites

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

Procedure

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

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

Category	Description
Edit Mode	<p>You can select the storage cost to be applicable for all the data centers or a specific data center.</p> <ul style="list-style-type: none"> ■ Edit for All Data Centers mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost. ■ Edit for specific Data Center mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.
Select Data center	You can select the data center for which you want to change the storage cost. This field is applicable only for specific data centers.
Tag Category	<ul style="list-style-type: none"> ■ Category displays the tag categories for datastores and also the tags associated with the category.
Datastores	Displays the total number of datastores for a specific category or type. You can click the datastore value to see the list of datastores and its details such as monthly cost, total GB for each datastore.
Total Storage (GB)	Displays the total storage for a specific category or type.
Monthly Cost Per GB	Displays the monthly cost per GB for a specific category or type. You can edit this value for defining the monthly cost per GB for datastores.
Monthly Cost	Displays the total monthly cost for a specific category or type.

4 Click **Save**.

Edit Monthly Cost of License

You can edit the total operating system licensing cost and VMware license cost of your cloud environment. You can now set a total fixed cost for the license in vRealize Operations Manager . The total license cost is divided across all the hosts present in the data center. You can edit the license cost by either selecting the ELA charging policy or selecting the per socket value.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Drivers**.
- 2 In the Cost Drivers tab, click **License**.
- 3 Select the required edit mode for changing the license cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
 - **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

4 Click **Save**.

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

Category	Description
Name	<p>Displays the category of the operating system. If the operating system is not Windows or Linux, cost drivers categorize the operating system under Other Operating Systems.</p> <hr/> <p>Note Two new cost components, Monthly cost of VMware vSAN Per Socket and Monthly cost of VMware vSAN SnS have been included for the vSAN cost calculation. The default values for these components are based on the reference database values.</p> <hr/> <p>The licensing cost for the Windows operating system falls under one of the following categories:</p> <p>Per Core License, applicable for</p> <ul style="list-style-type: none"> ■ Windows Server 2016 ■ Windows Server 2019 <p>Per Socket License, applicable for</p> <ul style="list-style-type: none"> ■ Windows NT 4.0 ■ Windows Server 2003 ■ Windows Server 2008 ■ Windows Server 2012 <p>Per Instance License, applicable for</p> <ul style="list-style-type: none"> ■ Windows XP ■ Windows Vista ■ Windows 98 ■ Windows 95 ■ Windows 8 ■ Windows 7 ■ Windows 3.1 ■ Windows 2000 ■ Windows 10
VMs	Displays the number of virtual machines that are running on the specific operating system.
Sockets	Displays the number of sockets on which the specific operating system is running.
Charged by	<p>Displays whether a cost is charged by socket or ELA.</p> <hr/> <p>Note The Charged By column can be edited to mention that the cost is charged by socket, core, instance, or ELA.</p>
Total Cost	Displays the total cost of the specific operating system.

5 Click **Save**.

Results

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

Customizing License Assignment

You can customize the licensing cost associated with your host using the custom license assignment option. Based on your requirement you can add or delete different operating system licenses to your host. With the custom license assignment option, you can increase or decrease the licensing cost associated with your host.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Drivers**.
 - 2 In the Cost Drivers tab, click **License**.
 - 3 Select the required edit mode for changing the monthly license cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
 - **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.
-
- Note** When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu
-
- 4 To customize the license cost for a specific server, click **Customize License Assignment**.
 - 5 Select the host for which you want to customize the license cost and click **Assign**.
 - 6 From the drop-down menu, select the operation system and click **Ok**.
The new operating system is listed under the Current Assignment column.
 - 7 To remove an existing operating system from the host, under **Current Assignment** click X icon next to the operating system.
The license cost of the removed operating system is reduced from the total cost.
 - 8 Click **Save**.
 - 9 Navigate to the **Cost Calculation Status** tab and click **Run**.

Results

The license cost is updated for the host, the * sign next to the host indicates that the license cost for the host has changed.

Category	Description
Server	You can select the server for which you want to customize the license cost.
Current Assignment	Displays the current operating systems associated with the host.
Default Assignment	Displays the default operating systems associated with the host.
Filter	Filters the hosts based on the operating system type.
Reset	Resets the license cost of the host to the default value.

Edit Monthly Cost of Maintenance

You can edit the monthly cost of maintaining your cloud environment. Maintenance cost is categorized into hardware maintenance cost and operating system maintenance cost. Hardware

maintenance cost is calculated as a percentage of the purchase cost of servers. Operating system maintenance cost is calculated as a percentage of the Windows licensing costs. You can now specify a total fixed cost for maintenance in vRealize Operations Manager . The total maintenance cost is divided across all the hosts present in the data center.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Drivers tab, click **Maintenance**.
- 3 Select the required edit mode for changing the monthly maintenance cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
 - **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

- 4 To customize the maintenance cost for a specific server, click **Edit For Individual Servers**.
- 5 Click **+Add Cost Per Server**.
- 6 From the **Select Server's for customization** drop-down select the required server and click **Ok**.
- 7 Specify the Server Hardware Percentage and OS Percentage and click **Save**.

View the change in maintenance cost after you have run the cost calculation cycle.

Edit Monthly Cost of Labor

You can edit the monthly cost of labor for your cloud environment. You can set a total fixed cost for labor in vRealize Operations Manager . The total labor cost is divided across all the hosts present in the data center. The labor cost is combination of the total cost of the server administrator, virtual infrastructure administrator, and the operating system administrator.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Labor**.
- 3 Select the required edit mode for changing the monthly labor cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

- 4 Edit the monthly labor cost.
 - Edit the detailed cost of labor.
 - Edit the total monthly labor cost for servers, virtual infrastructure, and operating system.
- 5 To customize the labor cost for a specific server, click **Server** and then click **Edit For Individual Servers**.
- 6 Click **+Add Cost Per Server**.
- 7 From the **Select Server's for customization** drop-down select the required server and click **Ok**.
- 8 Specify the Monthly hours of labor per hour, Labor hourly rate, and click **Save**.

The monthly labor cost is displayed.

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

Results

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

Edit Monthly Cost of the Network

You can edit the monthly cost for each Network Interface Controller (NIC) type or can edit the total cost of all the networking expenses associated with the cloud. You can now set a total fixed cost for network resources in vRealize Operations Manager . The total network cost is divided across all the hosts present in the data center.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Network**.
- 3 Select the required edit mode for changing the monthly network cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

Note When you select Edit for specific data center as the edit mode, then the select data center option is enabled. Select the data center from the drop-down menu

- 4 Edit the monthly cost of network.
 - Modify the values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and the 100 Gigabit NIC.
 - Modify the total monthly cost of all network expenses associated with the cloud.
- 5 To customize the network cost for a specific server, click **Edit For Individual Servers**.
- 6 Click **+Add Cost Per Server**.
- 7 From the **Select Server's for customization** drop-down select the required server and click **Ok**.
- 8 Specify values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and 100 Gigabit NIC and click **Save**.

View the change in network cost after you have run the cost calculation cycle.

Edit Monthly Cost of Facilities

For your cloud environment, you can specify the total monthly cost of facilities or edit the facilities cost for real estate, power, and cooling requirements. You can now set the total fixed cost for facilities in vRealize Operations Manager . The total facilities cost is divided across all the hosts present in the data center.

Procedure

- 1 Click **Administration** and in the left pane click **Configuration > Cost Settings**.
- 2 In the Cost Driver tab, click **Facilities**.
- 3 Select the required edit mode for changing the monthly facilities cost.
 - **Edit for All Data centers** - mode enables you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
 - **Edit for specific Data Center** - mode enables you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.
- 4 (Optional) Select the data center from the drop-down menu.

Note If you select Edit for specific data center as the edit mode, then the select data center option is enabled.

- 5 Edit the monthly facilities cost.
 - Modify the cost of rent or real estate per rack unit and modify the monthly cost of power and cooling per kilowatt-hour.

- Modify the total monthly cost of facilities.
- 6 To customize the facilities cost for a specific server, click **Edit For Individual Servers**.
 - 7 Click **+Add Cost Per Server**.
 - 8 From the **Select Server's for customization** drop-down select the required server and click **Ok**.
 - 9 Specify the Cost Per Kilowatt and Real Estate Cost Per Rack Unit and click **Save**.

View the change in network cost after you have run the cost calculation cycle.

Editing Additional Costs

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

Procedure

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

Note Additional cost driver allows you to assign costs at Host, vCenter, VM, cluster, or data center level. For example, if you want to keep a cluster protected using the disaster recovery services, which involves an additional cost of \$5000, you can do that by editing the additional cost driver.

- 4 Select the **Entity Type** and **Entity Selection**.

The **Entity Count** gets updated.

- 5 Enter the **Monthly Cost per entity** .

The **Total Cost per month** gets computed automatically.

- 6 Click **Save**.

Note After you update the Additional Cost configuration, you must reload the page manually to view the updated values.

Edit Application Cost

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

Prerequisites

Create applications in vRealize Operations Manager .

Procedure

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

Note You can now specify the cost of packaged applications that are discovered by the Service Discovery Management Pack. Earlier the option to specify the application cost was available only for business applications defined by the user.

- 4 Modify the cost of the application.
- 5 Click **Save**.

Cluster Cost Overview

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

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

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

Cluster Cost Elements	Calculation
Average Memory Utilization	Average Memory Utilization = (Cost attributed to Memory utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster).
Expected CPU Utilization	The utilization percentage level of CPU that the cluster is expected to operate. Note When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value.
Expected Memory Utilization	The utilization percentage level of Memory that the cluster is expected to operate. Note When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value.

Cluster Cost Computation with Allocation Model

You can now use the allocation model to compute the cost of clusters in vRealize Operations Manager, earlier the cluster cost computation was based on the cluster utilization. When you perform cost computation using the allocation model, you can set the over commit ratio for CPU, RAM, and storage.

Note The allocation ratio can be set at both cluster level and datastore cluster level. You can also mention the storage base rate, which will be displayed at the datastore level.

Table 4-61. Cluster Base Rate Computation with Allocation Model

Base Rate	Formula
vCPU Base Rate	vCPU base rate = B1 = (Cost attributed to CPU) / (Number of vCPUs in a cluster)
RAM Base Rate	RAM base rate = B2 = (Cost attributed to RAM) / Number of vRAMs in a cluster Note The cost computation is based on Over Commit ratio. If the Over Commit ratio is 1:4, and total cores in cluster are 6, then vCPU count = 24, in case if the allocated vCPU exceeds this targeted number, then the maximum value is selected.

Table 4-62. Virtual Machine Cost Computation with Allocation Model

Cost	Formula
Virtual Machine Cost	Virtual machine cost = (Number of vCPU allocated x B1 of cluster it belongs to) + Number of vRAMs allocated x B2 of cluster it belongs to) + storage cost + direct cost. Note Storage allocated represents the Storage Base Rate based on allocation.

Editing Cluster Cost Calculation Methods

You can edit the cluster cost calculation method based on your business requirement. The cost of a cluster is derived from cost drivers. Virtual machine cost is calculated by multiplying base rates with the utilization of the VMs.

Procedure

- 1 In the menu, Click **Administration** and then in the left pane click **Configuration > Cost Settings**.

- 2 In the Cluster Cost tab, click **CHANGE**.

The Cluster Cost Calculation Methods dialog box is displayed.

- 3 Select any one of the Cluster Cost Calculation methods.

Option	Description
Cluster Usable Capacity After HA and Buffer	<p>The cluster cost calculated total capacity minus resources needed for High Availability (HA) and the capacity buffer setting.</p> <p>Base rates are calculated based on the total cost of the cluster and Usable Capacity after HA and Buffer. Virtual machine costs are calculated from these base rates. Things to note:</p> <ul style="list-style-type: none"> ■ A lower buffer reduces the base rates and causes the virtual machines to become cheaper. ■ A higher buffer increases base rates and causes the virtual machines to become more expensive. ■ Base rates and virtual machine costs do not change with the utilization of the cluster. ■ The difference between Usable Capacity after HA and Buffer and actual utilization is used to compute unallocated costs.
Cluster Actual Utilization	<p>To calculate the base rates using the month to date average utilization of the cluster resources, select this option.</p> <p>Base rates are calculated based on the total cost of the cluster and average utilization. Virtual machine costs are calculated from these base rates. Things to note:</p> <ul style="list-style-type: none"> ■ Lower utilization level causes base rates to be high and virtual machines also become more expensive. ■ Higher utilization level causes base rates to be lower and virtual machines to become cheaper. ■ Base rates and virtual machine costs can change frequently based on the utilization of the cluster. ■ Unallocated cost of the cluster is near to zero. ■ The costs for unused resources are distributed across all virtual machines based on their actual utilization within the cluster.

- 4 Click **SAVE**.

Publish Daily Cost Metrics for Virtual Machines

In vRealize Operations Manager , you can now publish daily cost metrics for all virtual machines. The daily cost metric of a virtual machine is the sum of daily cost of CPU, memory, storage, and

additional cost associated with the virtual machine. Daily cost metrics provide granular details of the costs associated with the virtual machine.

Formula to Calculate the Daily Cost and Monthly Cost of Virtual Machines

You can calculate the daily cost associated with a virtual machine using the following formula.

Virtual Machine Cost Elements	Calculation
Daily Total Cost of Virtual Machine	Daily total cost of virtual machine = Sum of Daily cost of (CPU + memory + storage + additional cost)

The change in daily cost metrics also changes the way you calculate the effective month to date cost of a virtual machine. You can use the following formula to calculate Effective Month to Date cost for a virtual machine.

Virtual Machine Cost Elements for a Month	Calculation
Effective MTD Cost of VM	Sum of CPU daily cost from the beginning of the month until now + Sum of memory daily cost from the beginning of the month until now + Sum of storage daily cost from the beginning of the month until now + Sum of additional daily cost from the beginning of the month until now

How to View the Daily Cost Metrics of a Virtual Machine

To view the daily cost metrics of a virtual machine, from the menu, select **Administrator** and then in the left pane select **Inventory > vCenter Adapter**, select the specific **Virtual Machine**, and click the **Metrics** tab.

Pricing Overview

You can create pricing cards in vRealize Operations Manager to calculate the price associated with your virtual infrastructure. You can assign pricing cards to vCenters or Clusters, depending on the pricing strategy determined by vRealize Operations Manager administrator. The pricing cards help you to set the price for each resource present in your virtual environment.

You can customize the pricing card as per your requirement. vROps has two types of pricing cards, rate-based pricing card and cost-based pricing card. After configuring a pricing card, you can assign it to one or more vCenters or Clusters as determined by the pricing strategy.

How Is Price Calculated

In rate-based pricing policy vRealize Operations Manager calculates the virtual infrastructure price based on the rate card defined by you. For rate-based pricing policy vRealize Operations Manager lets you define cost elements as per your requirements.

The server recalculates the price every 24-hours, the price calculation for the new pricing cards is done in the next vRealize Operations Manager price calculation cycle.

Hierarchy of Pricing Policy

The assignment of policy in vRealize Operations Manager will be for Clusters and vCenters. The price is calculated for virtual machines, then it is aggregated and rolled up to vCenter. If there are two policies, a default policy for vCenter and another policy for Cluster, then the price calculation is based on the cluster policy for all the resources under the cluster. After that the cluster cost is rolled up to vCenter.

When a virtual machine is under vRealize Automation hierarchy and vCenter hierarchy, then the pricing is calculated based on the vRealize Automation hierarchy and the virtual machine is removed from the vCenter resources and included under vRealize Automation resources.

Pricing Support for VMware Cloud on AWS Resources

You can create a pricing policy in vRealize Operations Manager and assign it to VMware Cloud on AWS (VMC) resources, however you can only use the rate-based pricing policy for VMC-related objects.

Note When you assign Cost-Based Policy for VMC resources, the policy is not applied, and price calculated for the policy is reported as zero.

Add New Pricing Card

You can add and assign new pricing card to vCenter and Clusters in vRealize Operations Manager . The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

Procedure

- 1 Navigate to **Administration > Cost Settings > Pricing**.

- 2 Click **New Pricing Card** and configure the details of the pricing card.

Table 4-63. Pricing Card Configuration

Parameter	Description
Name and Description	<ol style="list-style-type: none"> 1 Enter a name and description for your pricing card. 2 Optional: Select Default for Unassigned Workloads. 3 Click Next. <p>Default pricing card applies to all vCenter resources which do not have a direct cost policy assigned to them.</p>
Basic Charges	<p>Select the type of pricing card. Follow the steps for cost-based pricing card.</p> <ol style="list-style-type: none"> 1 Enter the Cost Factor for the following. <ol style="list-style-type: none"> a CPU Cost b Memory Cost c Storage Cost d Additional Cost 2 Select the charging period as per your requirement, the options are Hourly, Daily, Weekly, and Monthly. 3 Select how to charge for the resources, the options are Always or Only When Powered On. 4 Click Next. <hr/> <p>Note Cost - The cost is defined in vRealize Operations. If selected, a multiplication factor is required. For example, if you select 1.1 as a factor, the cost is multiplied by 1.1 resulting in a 10% increase to the calculated cost. The price equation using cost is: <cost> x <multiplication factor> = Price</p> <p>Follow the steps for rate-based pricing card.</p> <ol style="list-style-type: none"> 1 Enter the CPU Rate in MHz per vCPU. 2 Enter the Memory Rate per GB. 3 Enter Storate Rate per GB. 4 Select the ChargingPeriod for all the values. 5 Select the Charge On Power State for all the values.
Guest OSes	<ol style="list-style-type: none"> 1 Enter the Guest OS Name. 2 Enter the base rate. 3 Select the charging period as per your requirement, the options are Hourly, Daily, Weekly, and Monthly.

Table 4-63. Pricing Card Configuration (continued)

Parameter	Description
Tags	<p>Enter the Tag name and Tag Value. Define the charging method and base rate.</p> <ul style="list-style-type: none"> ■ Recurring - enter a base rate and define recurring interval as the charge period. The absolute rate value is required and it is added to the overall price. ■ One time - define the one-time base rate charge. The absolute value is required and it is added as a one time price. ■ Rate Factor - A multiplication factor is required that is applied to the select charge category. <p>Select how to charge the Tag based on powered on state.</p>
Overall Charges	<p>You can define overall charges to VMs that match this policy.</p> <ol style="list-style-type: none"> 1 Enter the VM setup charges. 2 Enter the Recurring charge and select the time period from the drop-down menu.
Assignments	<p>You can assign the new pricing card to vCenters and Clusters.</p> <ol style="list-style-type: none"> 1 Select the vCenter or Cluster to which you want to apply the pricing card. 2 Click Add and Click Finish.

Results

The new pricing card details are displayed in the Pricing tab.

Cost Calculation Status Overview

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

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

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

vRealize Business for Cloud supports migration of cost driver configuration from vRealize Business for Cloud to vRealize Operations Manager . You can migrate cost driver configuration from vRealize Business for Cloud 7.x or later to vRealize Operations Manager 6.7 or vRealize Operations Manager 7.5.

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

Costing Enhancements

In vRealize Operations Manager , a new global property Cluster Utilization Ceiling Factor is introduced. Using Cluster Utilization Ceiling Factor, you can specify the ceiling value and calculate the base rate for a cluster.

You can use the ceiling factor only if the base rate cost calculation is done using Cluster Actual Utilization method. After you set the ceiling factor value, the Actual Utilization of the cluster is rounded off to the next available multiple of the ceiling value. When ceiling value is 0, Expected Utilization is equal to actual utilization. When ceiling value is 20, it is not considered as special case, actual utilization is rounded off to the next multiple.

Note The ceiling value range is from 0 to 20. If the number is out of this range, the default value of five is used as the ceiling number.

How to Set the Cluster Base Rate Calculation Method

To change the Cluster Base Rate Calculation method, you must go to **Administration > Configuration > Cost Settings > Cluster Cost** page. Click **Change** next to the Cluster Base Rate calculation method and select Cluster Actual Utilization.

Where to Find Cluster Utilization Ceiling Factor

To set the ceiling value for a cluster, you must go to **Administration > Management > Global Settings > Cluster Utilization Ceiling Factor**. Enter the ceiling value between 0 and 20 and click **Save**.

To view the change in cost metrics, run the Cost Calculation Status and select a cluster .

If the Actual Utilization of the cluster for CPU is 30 % and Memory is 45%, and the ceiling value specified is 10, then

- Cluster Expected CPU Utilization (%) = 40
- Cluster Memory Expected Utilization (%) = 50

Actual Cluster Utilization is rounded off to the ceiling value.

If you set the Cluster Utilization Ceiling Factor to either 0 or 20, then the value of Expected Memory Utilization changes to the next number. For example, if you set the ceiling factor to 0 then, the expected utilization value changes to 1.

Support to Roll up Name Space Cost Metrics

The cost metrics of Point of Delivery (Pod) virtual machines (VMs) has been enhanced to support the following scenarios:

- Cost metrics of Pod VMs are rolled up to the Name Space and Guest Cluster level.

- All the cost metrics of VMs, Pods, and guest cluster which are present under Name Space are rolled up to Name Space and Guest Cluster level.

Old Cost Metrics	Rolled up Cost Metrics
Effective MTD Total Cost	Aggregate Additional Daily Cost
Deleted VM Daily Cost	Aggregate Deleted VM Daily Cost
Daily CPU Cost	Aggregate CPU Daily Cost
Daily Memory Cost	Aggregate Memory Daily Cost
Daily Storage Cost	Aggregate Storage Daily Cost
Daily Additional Cost	Aggregate Additiona Daily Cost

vRealize Automation 8.X

The vRealize Automation 8.x extends operational management capabilities of the vRealize Operations Manager platform to provide the cloud aware operational visibility of the cloud infrastructure. The vRealize Automation 8.x enables you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use the vRealize Automation 8.x to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with vRealize Operations Manager .
- Import and synchronize existing cloud accounts from vRealize Automation 8.x to vRealize Operations Manager .
- Manage the workload placement of VMs that are part of the clusters managed by vRealize Automation 8.x.
- Integrate and troubleshoot vSphere endpoint issues associated with vRealize Automation 8.x using the vRealize Operations Manager dashboard.

Note In this release we support only vSphere endpoints.

vRealize Operations Manager and vRealize Automation Integration - Technical Overview

The vRealize Automation 8.x integration with vRealize Operations Manager , extends operational management capabilities of the vRealize Operations Manager platform to provide cloud aware operational visibility of the cloud infrastructure. The vRealize Automation 8.x enables you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use the vRealize Automation 8.x to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with vRealize Operations Manager .

- Import and synchronize existing cloud accounts from vRealize Automation 8.x to vRealize Operations Manager .
- Manage the workload placement of VMs that are part of the clusters managed by vRealize Automation 8.x.
- Integrate and troubleshoot vSphere endpoint issues associated with vRealize Automation 8.x using the vRealize Operations Manager dashboard.

How Does vRealize Automation and vRealize Operations Manager Integration Work

vRealize Automation can work with vRealize Operations Manager to perform advanced workload placement, provide deployment health and virtual machine metrics, and display pricing.

Integration between the two products must be on-premises to on-premises, not a mix of on-premises and cloud.

To integrate with vRealize Operations Manager , look under **Infrastructure > Connections > Integrations**. To add the integration, you need the vRealize Operations Manager URL and its login user name and password. In addition, vRealize Automation and vRealize Operations Manager need to manage the same endpoint.

Workload Placement

When you deploy a blueprint, workload placement uses collected data to recommend where to deploy the blueprint based on available resources. vRealize Automation and vRealize Operations Manager work together to provide placement recommendations for workloads in the deployment of new blueprints.

While vRealize Automation manages organizational policies, such as business groups, reservations, and quotas, it integrates with the capacity analytics of vRealize Operations Manager to place machines. Workload placement is only available for vSphere endpoints.

Workload Placement Terms Used

Several terms are used with workload placement.

- Clusters in vSphere map to compute resources in vRealize Automation.
- Reservations include compute and storage, where the storage can consist of individual datastores or datastore clusters. A reservation can include multiple datastores, datastore clusters, or both.
- Multiple reservations can refer to the same cluster.
- Virtual machines can move to multiple clusters.
- When workload placement is enabled, the provisioning workflow uses the placement policy to recommend where to deploy the blueprint.

Provisioning Blueprints with Workload Placement

When you use workload placement to provision blueprints, the provisioning workflow uses the reservations in vRealize Automation, and the placement optimization from vRealize Operations Manager .

- 1 vRealize Operations Manager provides placement optimization recommendations according to analytics data.
- 2 vRealize Automation continues the provisioning process according to the placement recommendations from vRealize Operations Manager .

If vRealize Operations Manager cannot provide a recommendation, or the recommendation cannot be used, then vRealize Automation falls back to its default placement logic.

Workload Placement Goals

The goal of Workload Placement is to make sure that no cluster is overloaded by more than 80% of the potential workload. Workload placement is done in the following three stages.

Stress-free clusters

Ensures that the memory, CPU, or disk space workload is less than 80% for the cluster.

Workload Placement based on Business Intent

Distribution of virtual machines between Clusters is based on tags. When a cluster and VM have the same tag, VM will be recommended to move from this cluster or VM will be recommended to move to this cluster. When host-based tagging is enabled, VM will be recommended to optimize the workload for cluster based on a rule.

Distribution Strategy

- Balanced distribution: Distribution is based on the green zone, with maximum 20% difference workload between the two clusters.
- Moderate distribution: Ensures that no cluster is at stress level.
- Consolidated distribution: Keeps the hosts free while maintaining the workload at green level. In some cases, one of the clusters has resources free for backup purposes.

Workload Placement Recommendation

Workload Placement is recommended to run on a cluster (with existing VMs) or for the new deployment in vRealize Automation for Day 0 integration. After the deployment or move of the virtual machine, the cluster hosting that VM does not have workload greater than 80% for CPU and/or Memory and/or Disk Space. The recommendation starts only if the Memory or CPU workload is not optimized.

Note We do not recommend Disk Space Optimization for Workload Placement, as we always ensure that the workload for Disk Space is within the green zone.

vRealize Automation Workload Placement Day 1 Recommendation

The distribution of VMs is done based on blueprint configurations. WLP calculates and evaluates the impact of potential deployment based on the workload or cluster utilization. The objective of WLP is to make sure that the least loaded cluster gets to provision highest number of VMs.

We have cluster A which has 100 GB memory capacity of which 20 GB is free, so this means 80 GB is used. We have another cluster B which has 1 TB memory of which 700 GB is free, so this means that 300 GB is used. If you look at this percentage wise, we see that cluster A has 80% free space and cluster B has 70% free space, however in terms of actual available space, we see that the 700 GB free space of cluster B is more than 20 GB free space available in cluster A.

Note If the workload placement results in having more than 80% workload on the cluster, then vRealize Operations Manager cannot provide a recommendation, or the recommendation cannot be used, then vRealize Automation falls back to its default placement logic.

Workload Placement Automation

Automation

The Automation calculates and evaluates the move of virtual machines every 5 minutes. If you find a VM that is not optimized, the optimization is triggered automatically. Note, the time slot between two automated optimizations is limited to 6 hours.

Schedule

Schedule automation calculates and evaluates the move only during the scheduled time slots. The available options are Once, Daily, Weekly, and Monthly.

Impact on Cloud Zones and Non vRealize Automation-Managed VMs

Whenever there is a vRealize Operations and vRealize Automation integration for a data center, the cloud zones that have virtual machines which are not managed by vRealize Automation or not created by vRealize Automation, Workload Placement ignores them.

Supported vRealize Automation Versions

vRealize Automation 8.x is supported on vRealize Operations Manager 8.3 version. Workload placement for day 1 operations is supported from vRealize Automation 7.3 onwards with vRealize Operations Manager 6.6 and above. Workload placement for day 2 operations is supported from vRealize Automation 7.5 onwards with vRealize Operations Manager 7.0 and above.

Object Types

vRealize Automation 8.x brings in cloud accounts and their relationships from vRealize Automation into vRealize Operations Manager for operational analysis. You can use the following items in the virtual infrastructure as object types in vRealize Operations Manager .

- Cloud Zone
- Blueprint
- Project

- Deployment
- Cloud Account
- User
- Organization
- Cloud Automation Services World

Workload Placement

In vRealize Operations Manager , you can configure vRealize Automation 8.x instances to work with vRealize Operations Manager instances. Using vRealize Operations Manager you can monitor the placement of existing workloads and optimize the resource usage.

Prerequisites

- Verify that the user has privileges of Organizational Owner and Cloud Assembly Administrator set in vRealize Automation.
- You must know the vCenter Server credentials and have the necessary permissions to connect and collect data.
- Verify that vRealize Automation 8.x is enabled from **Administration > Management > Integrations** in vRealize Operations Manager . For more information, see [Configuring vRealize Automation 8.x with vRealize Operations Manager](#).
- vRealize Operations Manager must have the same vCenter Cloud Account configured to match with vRealize Automation 8.x.
- Ensure that integration is enabled for vRealize Operations Manager and vRealize Automation 8.x.

Procedure

- 1 In the menu, select **Home** and then select **Workload Optimization**.
- 2 Click the **View** filter drop-down menu and select the **VRA Managed** objects.
All the Cloud Zones related to the vCenter Server are displayed in vRealize Operations Manager .
- 3 Click the **Cloud Zone** you want to optimize.
- 4 Based on the operational intent, click **Optimize Now**.
The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.
- 5 If you are satisfied with the projected results of the optimization action, click **NEXT**.

6 Review the optimization moves, then click **BEGIN ACTION**.

In the scope of vRealize Automation 8.x integration, vRealize Operations Manager sends a move migration request directly to vRealize Automation 8.x. In the earlier versions, the migration request was sent to the vCenter Server.

What to do next

To verify that the optimization action is complete, select **Administration** on the top menu, and click **History > Recent Tasks** in the left pane. In the **Recent Tasks** page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

Pricing for vRealize Automation 8.x Components in vRealize Operations Manager

After you integrate vRealize Automation 8.x private cloud adapter instances with vRealize Operations Manager, you can calculate the cost of deployments, projects, and virtual machines of the selected cloud adapter. Pricing provides an overview of the costs related to the cloud environment, cloud resources, and the costs associated with the project.

How the Pricing Works in vRealize Automation 8.x

- vRealize Operations Manager understands the constructs defined in vRealize Automation 8.x and calculates the CPU, RAM, Storage and Additional prices for Projects, Deployments, and virtual machines.
- A single project can have multiple deployments and a single deployment can have multiple virtual machines associated with the deployment.
- Pricing for multiple virtual machines associated with the deployment is the sum of all the resources associated with individual virtual machines.
- If a single project has multiple deployments, then the project pricing is equal to the sum of individual deployments. The deployment can have multiple virtual machines and resources associated with it.
- On day one, the pricing is equal to the cost of resources defined in vRealize Operations Manager.
- On day two, the price is calculated using the following formula.
 - Cost of resources for the present day – Cost of resources for the previous day
- If in case the pricing does not happen as per the definition, then the partial price is set to true, and the pricing is calculated based on the previous days price.
- In vRealize Operations Manager, the following new dashboards are included to view the pricing details for the vRealize Automation 8.x instances.
 - Cloud Automation Environment Overview

- Cloud Automation Project Cost Overview
- Cloud Automation Resource Consumption Overview
- Cloud Automation Top-N Dashboard

Data Collection Enhancements in vRealize Automation for Pricing in vRealize Operations Manager

The following enhancements have been made for the data collection process from vRealize Automation for pricing purposes.

- Collect cloud zones with relation to clusters and resources pools from vRealize Automation to vRealize Operations Manager .
- Collect Projects from vRealize Automation with relation to deployments.
- Include project, cloud zone, and blueprint as properties in virtual machines that are deployed in vRealize Automation.

Upfront Price Support for vRealize Automation 8.x Private Cloud Components

vRealize Operations Manager supports upfront pricing for vRealize Automation 8.x in the following ways:

- vRealize Operations Manager uses rate cards to provide upfront cost estimates of catalog items just before deployment.
- vRealize Automation 8.x retrieves the deployment cost and estimated cost from vRealize Operations Manager .
- vRealize Automation user interface allows you to customize the pricing policies and assign them to the projects or cloud zones.
- If vRealize Automation does not specify the pricing policy, then the price is calculated using the vRealize Operations Manager cost calculation policy.
- If a custom pricing policy is set for a price calculation, then the deployment and upfront catalog price computation is done as per the custom policy.

Upfront Price Support for VMware Cloud on AWS Resources

vRealize Operations Manager supports upfront pricing for VMware Cloud on AWS resources in the following ways:

- vRealize Operations Manager supports upfront pricing for VMware Cloud on AWS only if rate-based pricing is configured in vRealize Automation for VMware Cloud on AWS resources.
- vRealize Operations Manager does not support cost-based computation for VMware Cloud on AWS resources.

Configuring vRealize Automation 8.x with vRealize Operations Manager

To access vRealize Automation 8.x instance and troubleshoot automation issues using vRealize Operations Manager , you must configure the vRealize Automation adapter in vRealize Operations Manager .

Prerequisites

- Verify that you know the FQDN/IP address, user name, and password of the vRealize Automation instance you have installed.
- Ensure that the vRealize Automation user has both organizational owner and Cloud Automation Services administrator permissions.

Procedure

- 1 In the menu, select **Administration** and then from the left pane, select **Management > Integrations**.
- 2 From the **Integrations** page, click vRealize Automation8.x.
- 3 In the **vRealize Automation 8.x** page, enter the FQDN or IP address of the vRealize Automation 8.x instance to which you want to connect.
- 4 Set **Auto Discovery** to true.
- 5 To add credentials, click the plus sign.
 - a In the Credential name text box, enter the name by which you are identifying the configured credentials.
 - b Enter the user name and password for the VMware vRealize Automation instance.
 - c Click **OK**.

You have configured credentials to connect to a VMware vRealize Automation instance.
- 6 From the **Collectors/Groups** drop-down menu, select the collector group.
- 7 Click **Validate Connection** to verify that the connection is successful.
- 8 Review and accept the Server certificate.
- 9 Click **Advanced Settings**.
- 10 From the **User Count** drop-down menu, select the number of user resources to be imported from vRealize Automation.

The User Count options are 20, 100, 200, 300, 400, and All Users.
- 11 Click **Save** to save the adapter instance.

Results

After integrating vRealize Automation adapter instance with vRealize Operations Manager , you can view the vRealize Automation adapter data from the vRealize Operations Manager dashboard.

Cloud Zones in vRealize Operations Manager

Cloud zones enable you to group a set of compute resources and assign capability tags to the zone. The cloud zone is based on accounts/regions, so you must have at least one cloud account configured before you can create a cloud zone. Cloud zones define where and how blueprints configure deployments. You can have one or many cloud zones assigned to each project based on priority and limits.

How Cloud Zones Work

After you integrate vRealize Automation 8.x with vRealize Operations Manager , you can retrieve cloud zones into vRealize Operations Manager . The **Cloud Zones** option is hidden from the user until the integration with vRealize Automation 8.x is enabled from the integration page under **Administration > Management**.

The Cloud Zones option is enabled in vRealize Operations Manager , only if the following conditions are met.

- vRealize Automation 8.x instance is integrated successfully in vRealize Operations Manager **Administration > Management>Integrations**.
- vRealize Automation 8.x objects are discovered in vRealize Operations Manager .
- vRealize Automation 8.x accounts and vRealize Operations vCenter Cloud Accounts are synchronized.

All the Cloud Zone objects which are existing in vRealize Automation 8.x environment, are discovered in vRealize Operations Manager . Cloud zones, whose dependent clusters are not discovered in vRealize Operations Manager , are not represented in Capacity Overview, Reclaim, and Workload Optimization pages.

Cloud Zones List

You can view the list of cloud zones that exist in your environment. In this view, you can click a cloud zone to display all the resources and objects that are associated with the cloud account. When you click the Cloud Zone, you are directed to the standard object summary page of the cloud account.

Where You Find Cloud Zones

Select **Environment** in the menu and click **Cloud Zones** tab.

Cloud Zone Tab Options

Option	Description
Name	Displays the name of the selected cloud zone.
Cloud Account	Displays the cloud accounts associated with the cloud zone.

Option	Description
Resources	<p>Displays the cloud account resources associated with the cloud zone.</p> <hr/> <p>Note If the resource field is empty, it means vRealize Operations Manager does not have a corresponding vCenter Cloud Account for that associated Cloud Zone. Add a new vCenter Cloud Account manually or use the Import Cloud Account option from the Cloud Account page.</p> <hr/>
Capability Tags	Displays the capability tags associated with the cloud zone.

vSAN

You can make vSAN operational in a production environment by using dashboards to evaluate, manage, and optimize the performance of vSAN objects and vSAN-enabled objects in your vCenter Server system.

vSAN extends the following features:

- Discovers vSAN disk groups in a vSAN datastore.
- Identifies the vSAN-enabled cluster compute resource, host system, and datastore objects in a vCenter Server system.
- Automatically adds related vCenter Server components that are in the monitoring state.
- Support for vSAN datastores in workload optimization with cross-cluster rebalance actions.
 - You can move VMs from one vSAN datastore to another vSAN datastore.
 - You can optimize the container if all the vSAN clusters are not in resync state.
 - VMs with different storage policies for each disk or VMs with different types of storage for each disk will not be moved.
 - You can generate a rebalance plan only if sufficient disk space is available at the destination vSAN datastore (The vSAN datastore slack space will also be considered).
 - The storage policy assigned to the VM will be considered during the workload optimization (Compatibility check is performed against the storage policy).
 - VM migration from vSAN datastore to vSAN stretched clusters is not supported.

Configure a vSAN Adapter Instance

When configuring an adapter instance for vSAN, you add credentials for a vCenter Server. In the earlier versions of vRealize Operations Manager, the vSAN solution was installed as part of the vRealize Operations Manager installation. Now, in case of a new installation the vSAN solution is pre-bundled as part of vRealize Operations Manager OVF, you must install the vSAN solution separately.

Prerequisites

Only vCenter Server systems that are configured for both the vCenter adapter and the vSAN adapter appear in the inventory tree under the vSAN and Storage Devices. Verify that the vCenter Server that you use to configure the vSAN adapter instance is also configured as a vCenter adapter instance for the VMware vSphere® solution. If not, add a vCenter adapter instance for that vCenter Server.

You must open port 5989 between the host and any vRealize Operations Manager node on which the vSAN adapter resides. This is applicable when the vSAN version in vSphere is 6.6 or lower.

You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to monitor the vSAN and Storage Devices.

To know how to install the Native Management Packs, see [Solutions Repository](#).

Procedure

- 1 In the menu, select **Administration** and then select **Solutions > Cloud Accounts** from the left panel.
- 2 From the **Cloud Accounts** page, select the vCenter Server instance from the list and then click the **vSAN** tab.
- 3 To use the vCenter Server for enabling vSAN, move the **vSAN configuraton** option to the right.

Note Once vSAN adapter instance is enabled and saved, the enable vSAN configuration option is not visible.

- 4 The credentials provided for the vCenter Server instance are also used for vSAN adapter instance. If you do not want to use these credentials, you can click **Use alternate credentials** option.
 - a Click the plus sign next to the Credential field and enter the details in the **Manage Credentials** dialog box.
 - b Enter the credential name, vCenter user name, and password and click **OK**.
- 5 Choose **Enable SMART data collection**, to enable SMART data collection for physical disk devices.
- 6 Click **Add**.

The vSAN configuration is enabled for the cloud account.
- 7 Click **Test Connection** to validate the connection with your vCenter Server instance.
- 8 Accept the vCenter Server security certificate.
- 9 Click **Save Settings**.

Results

The adapter is added to the Adapter Instance list and is active.

What to do next

To verify that the adapter is configured and collecting data from vSAN objects, wait a few collection cycles, then view application-related data.

- **Inventory.** Verify that all the objects related to the vSAN instance are listed. Objects should be in the collecting state and receiving data.
- **Dashboards.** Verify that vSAN Capacity Overview, Migrate to vSAN, vSAN Operations Overview, and Troubleshoot vSAN, are added to the default dashboards.
- Under **Environment > vSAN and Storage Devices**, verify that the vSAN hierarchy includes the following related vCenter Server system objects:
 - vSAN World
 - Cache Disk
 - Capacity Disk
 - vSAN-enabled vCenter Server clusters
 - vSAN Fault Domains (optional)
 - vSAN-enabled Hosts
 - vSAN Datastores
 - vSAN Disk Groups
 - vSAN Datastore related VMs
 - vSAN Witness Hosts (optional)

Verify that the Adapter Instance is Connected and Collecting Data

You configured an adapter instance of vSAN with credentials for a vCenter Server. Now you want to verify that your adapter instance can retrieve information from vSAN objects in your environment.

To view the object types, in the menu, click **Administration > Configuration > Inventory > Adapter Instances > vSAN Adapter Instance > <User_Created_Instance>**.

Table 4-64. 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.

Table 4-64. Object Types that vSAN Discovers (continued)

Object Type	Description
vSAN Witness Host	A tag for a witness host of a stretched cluster, if the stretched cluster feature is enabled on the vSAN cluster.
vSAN World	A vSAN World is a group parent resource for all vSAN adapter instances. vSAN World displays aggregated data of all adapter instances and a single root object of the entire vSAN hierarchy.
Cache Disk	A local physical device on a host used for storing VM files in vSAN.
Capacity Disk	A local physical device on a host used for read or write caching in vSAN

The vSAN adapter also monitors the following objects discovered by the VMware vSphere adapter.

- Cluster Compute Resources
- Host System
- Datastore

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Configuration > Inventory**.
- 2 In the list of tags, expand **Adapter Instances** and expand **vSAN Adapter Instance**.
- 3 Select the adapter instance name to display the list of objects discovered by your adapter instance.
- 4 Slide the display bar to the right to view the object status.

Object Status	Description
Collection State	If green, the object is connected.
Collection Status	If green, the adapter is retrieving data from the object.

- 5 Deselect the adapter instance name and expand the **Object Types** tag.

Each Object Type name appears with the number of objects of that type in your environment.

What to do next

If objects are missing or not transmitting data, check to confirm that the object is connected. Then check for related alerts.

To ensure that the vSAN adapter can collect all performance data, the Virtual SAN performance service must be enabled in vSphere. For instructions on how to enable the service, see [Turn on Virtual SAN Performance Service in the VMware Virtual SAN documentation](#).

If the Virtual SAN performance service is disabled or experiencing issues, an alert is triggered for the vSAN adapter instance and the following errors appear in the adapter logs.

```
ERROR com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- Failed to collect performance metrics for Disk Group
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- vSAN Performance Service might be turned OFF.
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
- (vim.fault.NotFound)
{
  faultCause = null,
  faultMessage = (vmobl.LocalizableMessage)
    [
      com.vmware.vim.binding.impl.vmodl.LocalizableMessageImpl@98e1294
    ]
}
```

vSAN Log Analytics Enhancements

When vRealize Operations Manager is integrated with vRealize Log Insight, you can view and troubleshoot vRealize Log Insight object issues within vRealize Operations Manager. Earlier you could troubleshoot issues related only to vCenter objects, but now you can troubleshoot issues related to vSAN also.

The enhancements to vSAN log analytics include use of specific queries to retrieve log information for the following vSAN objects:

- vSAN Cluster
- Witness Host
- Disk Group
- Cache Disk
- Capacity Disk

Where You Find vSAN Object Logs

Navigate to the vSAN Object Details page, and click the **Logs** tab.

Note If you are not logged in to vRealize Log Insight, then vRealize Operations Manager prompts you to log in to vRealize Log Insight with your login credentials.

vRealize Operations Manager uses special queries for each object type. Using the special queries for vSAN objects, you can perform the following actions:

- View interactive analytics for the selected vSAN object.
- Retrieve log details for the vSAN object.
- Analyze and troubleshoot issues related to the vSAN object.

vRealize Network Insight

The vRealize Network Insight adapter enables integration of vRealize Operations Manager with vRealize Network Insight. VMware vRealize Network Insight provides network visibility and analytics to minimize risk during application migration, optimize network performance, manage and scale VMware NSX-T, VMware NSX for vSphere, vCenter on VMware Cloud on AWS, VMware SD-WAN by VeloCloud, and Kubernetes deployments.

This adapter gets problem events from vRealize Network Insight and publishes the alerts in vRealize Operations Manager. Alerts are mapped correctly to the common objects between vRealize Network Insight and vRealize Operations Manager. Common objects supported in this adapter are vCenter Server, VMware NSX-T, and VMware NSX for vSphere. For the common objects, vRealize Operations supports launch-in-context to vRealize Network Insight. This allows the user to perform deep network troubleshooting with the vRealize Network Insight as the context.

The vRealize Network Insight adapter only supports vRealize Network Insight versions 5.2 and above. The vRealize Network Insight adapter can be installed and configured with On-prem version of vRealize Operations Manager or cloud version of vRealize Operations Cloud. The vRealize Network Insight adapter does not support cross platform configuration, it should be On-prem vRealize Operations Manager to On-prem vRealize Network Insight and vRealize Operations Cloud to vRealize Network Insight Cloud.

Configuring vRealize Network Insight

Configure an instance of the vRealize Network Insight in vRealize Operations Manager.

Prerequisites

As vCenter and NSX-T are native vRealize Operations Manager Management Packs, ensure that you have installed the latest NSX for vSphere Management Pack if you have NSX for vSphere data source configured in vRealize Network Insight.

Procedure

- 1 On the menu, click **Administration**.
- 2 In the left pane, expand **Management** and click **Integrations**.
- 3 Under Integrations, click the verticle ellipse next to VMware vRealize Network Insight and click **Configure**.

4 Configure the adapter instance.

Option	Description
VRNI FQDN/IP	The FQDN or the IP address of vRealize Network Insight.
Credential	<p>Select and add the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the environment of this management pack, click the plus sign.</p> <ul style="list-style-type: none"> ■ Credential Kind. Select and configure the Credential Type. You can select either the Local, LDAP, or vIDM network insight credentials. <p>Note This Management pack supports only the Local, LDAP, and vIDM users that are added in the User Management settings of vRealize Network Insight.</p> <ul style="list-style-type: none"> ■ Local - Network Insight Credentials. Enter the credential name, user name of the local user configured in vRealize Network Insight, and password for that user. ■ LDAP - Network Insight Credentials. Enter the credential name, LDAP domain configured in vRealize Network Insight, LDAP user name, and LDAP password for that LDAP user. ■ vIDM - Network Insight Credentials. Enter the credential name, vIDM FQDN/IP integrated with vRealize Network Insight, vIDM user name, and vIDM password for that vIDM user. <p>Credential Name. Credential Name.</p>
Collector / Group	Select the required collector group.
Validate Connection	Test Connection should be successful.

- 5 The vRealize Network Insight instance collects events based on common data sources between vRealize Operations Manager and vRealize Network Insight. When you disable the **Import problem events as based on common data sources** option, all the events are imported into the vRealize Operations Manager.
- 6 You can collect user-defined events of vRealize Network Insight as notifications in vRealize Operations Manager. To do so, enable the **Import User defined events as Notifications**.
- 7 Select the severity of the problem events you want to import. By default, all the problem events with moderate and critical severities are imported.
- 8 Click **Add**.

The vRealize Network Insight instance is added to the list.

End Point Operations Management Solution in vRealize Operations Manager

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

End Point Operations Management Agent Installation and Deployment

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

Prepare to Install the End Point Operations Management Agent

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

Prerequisites

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

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

- Define the agent `HQ_JAVA_HOME` location.

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

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

Supported Operating Systems for the End Point Operations Management Agent

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

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

Table 4-65. 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

Table 4-65. Supported Operating Systems for the End Point Operations Management Agent (continued)

Operating System	Processor Architecture	JVM
AIX 6.1, 7.1	Power PC	IBM Java SE7
VMware Photon Linux 1.0	x86_64	Open JDK 1.8.0_72-BLFS
Oracle Linux versions 5, 6, 7	x86_64, x86_32	Open JDK Runtime Environment 1.7

Selecting an Agent Installer Package

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Install the Agent on a Linux Platform from an RPM Package

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

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

The RPM performs the following actions:

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

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

Prerequisites

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

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

- If you are installing a non-JRE package, define the agent `HQ_JAVA_HOME` location.
End Point Operations Management platform-specific installers include JRE 1.8.x. Platform-independent installers do not. Depending on your environment and the installer you use, you might need to define the location of the JRE to ensure that the agent can find the JRE to use. See [Configuring JRE Locations for End Point Operations Management Components](#).
- If you are installing a non-JRE package, verify that you are using the latest Java version. You might be exposed to security risks with earlier versions of Java.
- Verify that the installation directory for the End Point Operations Management agent does not contain a vRealize Hyperic agent installation.
- If you are using the `noarch` installation, verify that a JDK or JRE is installed on the platform.

- Verify that you use only ASCII characters when specifying the agent installation path. If you want to use non-ASCII characters, you must set the encoding of the Linux machine and SSH client application to UTF-8.

Procedure

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

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

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

Results

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

What to do next

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

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

Install the Agent on a Linux Platform from an Archive

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

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

Prerequisites

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

Procedure

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

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

- 2 Run `cd agent_name/bin` to open the `bin` directory for the agent.

- 3 Run `ep-agent.sh start`.

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

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

What to do next

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

Install the Agent on a Windows Platform from an Archive

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

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

Prerequisites

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

Procedure

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

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

- 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 4-66. Silent Command Line Installer Parameters

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

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

Table 4-67. Additional Silent Command Line Installer Parameters

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

Install the Agent on an AIX Platform

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

Prerequisites

- 1 Install IBM Java 7.
- 2 Add the latest JCE from the IBM JRE security directory: `JAVA_INSTALLATION_DIR/jre/lib/security`.

Procedure

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

Install the Agent on a Solaris Platform

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

Prerequisites

- 1 Install Java 7 or above for Solaris from the Oracle site: https://java.com/en/download/help/solaris_install.xml
- 2 Add the latest JCE from <http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>

Procedure

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

Java Prerequisites for the End Point Operations Management Agent

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

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

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

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

Configuring JRE Locations for End Point Operations Management Components

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

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

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

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

How the Agent Resolves its JRE

The agent resolves its JRE based on platform type.

UNIX-like Platforms

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

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

Linux Platforms

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

Windows Platforms

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

- 1 `HQ_JAVA_HOME` environment variable

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

- 2 Embedded JRE

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

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

System Prerequisites for the End Point Operations Management Agent

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

As a workaround, complete the following steps:

Procedure

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

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

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

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

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

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

A number of steps are required to complete the configuration.

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

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

Prerequisites

Verify that the vRealize Operations Manager server is running.

Procedure

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

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

2 Specify the End Point Operations Management Agent Setup Properties

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

3 Configure an End Point Operations Management Agent Keystore

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

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

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

5 Overriding Agent Configuration Properties

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

6 End Point Operations Management Agent Properties

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

What to do next

Start the End Point Operations Management agent.

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

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

Procedure

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

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

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

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

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

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

Specify the End Point Operations Management Agent Setup Properties

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

Agent-server setup requires a minimum set of properties.

Procedure

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

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

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

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

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

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

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

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

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

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

Configure an End Point Operations Management Agent Keystore

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

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

Procedure

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

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

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

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

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

The agent configuration dialog box appears in these cases:

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

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

Prerequisites

Verify that the server is running.

Procedure

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

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

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

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

Results

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

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

Overriding Agent Configuration Properties

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

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

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

End Point Operations Management Agent Properties

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

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

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

Encrypt End Point Operations Management Agent Property Values

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

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

```
agent.setup.serverPword = ENC(4FyUf6m/
c5i+RriaNpSEQ1WKGb4y+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 out `hqadmin`.

agent.setup.camPort Property

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

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

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

Default

Commented out `7080`.

agent.setup.camPword Property

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

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

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

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

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

Default

Commented out `hqadmin`.

agent.setup.camSecure

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

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

Use `no=unencrypted` for unencrypted communication.

agent.setup.camSSLPort Property

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

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

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

Default

Commented out `7443`.

agent.setup.resetupToken Property

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

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

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

Default

Commented out `no`.

`agent.setup.unidirectional` Property

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

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

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

Default

Commented out `no`.

`agent.startupTimeout` Property

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

Default

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

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

`autoinventory.defaultScan.interval.millis` Property

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

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

Default

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

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

`autoinventory.runtimeScan.interval.millis` Property

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

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

Default

`86,400,000` milliseconds, or one day.

`http.useragent` Property

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

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

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

Default

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

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

log4j Properties

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

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

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

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

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

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

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

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

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

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

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



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

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

platform.log_track.eventfmt Property

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

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

Default

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

Attribute	Description
Timestamp	When the event occurred
Log Message	A text string
EventLogName	The Windows event log type System, Security, Or Application
EventAttributes	A colon delimited string made of the Windows event Source and Message attributes

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

Configuration

Use the following parameters to configure the Windows event attributes that the agent writes for a Windows event. Each parameter maps to Windows event attribute of the same name.

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

For example, with the property setting `platform.log_track.eventfmt=%user%%computer%
%source%:%event%:%message%`, the End Point Operations Management agent writes the following data when logging the Windows event 04/19/2010 06:06 AM Log Message (SYSTEM): SYSTEM: HP_Administrator@Office Print:7:Printer HP LaserJet 6P was paused.. This entry is for a Windows event written to the Windows system event log at 6:06 AM on 04/19/2010. The software associated with the event was running as "HP_Administrator" on the host "Office". The Windows event's Source, Event, and Message attributes, are "Print", "7", and "Printer HP LaserJet 6P was paused.", respectively.

plugins.exclude Property

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

Usage

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

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

plugins.include Property

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

Usage

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

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

postgresql.database.name.format Property

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

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

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

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

```
Database ${db}
```

where

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

Default

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

postgresql.index.name.format Property

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

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

Variable	Description
<code>DatabaseName</code>	The auto-discovered name of the database.
<code>Schema</code>	The auto-discovered schema for the database.
<code>Index</code>	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
<code>db</code>	Identifies the platform that hosts the PostgreSQL or vPostgreSQL server.
<code>schema</code>	Identifies the schema associated with the table.
<code>index</code>	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
<code>Host</code>	The FQDN of the platform that hosts the server.
<code>Port</code>	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
<code>postgresql.host</code>	Identifies the FQDN of the hosting platform.
<code>postgresql.port</code>	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
<code>DatabaseName</code>	The auto-discovered name of the database.
<code>Schema</code>	The auto-discovered schema for the database.
<code>Table</code>	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
<code>db</code>	Identifies the platform that hosts the PostgreSQL or vPostgreSQL server.
<code>schema</code>	Identifies the schema associated with the table.
<code>table</code>	The table name in PostgreSQL.

Default

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

`scheduleThread.cancelTimeout` Property

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

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

Usage

```
scheduleThread.cancelTimeout=5000
```

Default

5000 milliseconds.

scheduleThread.fetchLogTimeout Property

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

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

Usage

```
scheduleThread.fetchLogTimeout=2000
```

Default

2000 milliseconds.

scheduleThread.poolsize Property

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

Usage

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

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

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

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

Default

1

scheduleThread.queueSize Property

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

Usage

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

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

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

For example,

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

Default

1000

sigar.mirror.procnets Property

mirror /proc/net/tcp on Linux.

Default

true

sigar.pdh.enableTranslation Property

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

snmpTrapReceiver.listenAddress Property

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

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

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

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

Usage

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

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

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

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

Managing Agent Registration on vRealize Operations Manager Servers

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

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

You must have AgentManager credentials to perform the agent registration process. On a freshly deployed instance of vRealize Operations Manager, before you register the End Point Operations Management agent, you must also manually activate the management pack from **Administration > Solutions > Repository > Operating Systems/Remote Service Monitoring**.

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

Regenerate an Agent Client Certificate

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

Prerequisites

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

Procedure

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

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

Results

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

Securing Communications with the Server

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

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

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

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

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

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

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

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

Either the SHA1 or SHA256 algorithm can be used for the thumbprint.

Launching Agents from a Command Line

You can launch agents from a command line on both Linux and Windows operating systems.

Use the appropriate process for your operating system.

If you are deleting the `data` directory, do not use Windows Services to stop and start an End Point Operations Management agent. Stop the agent using `epops-agent.bat stop`. Delete the `data` directory, then start the agent using `epops-agent.bat start`.

Run the Agent Launcher from a Linux Command Line

You can initiate the agent launcher and agent lifecycle commands with the `epops-agent.sh` script in the `AgentHome/bin` directory.

Procedure

- 1 Open a command shell or terminal window.
- 2 Enter the required command, using the format `sh epops-agent.sh command`, where `command` is one of the following.

Option	Description
<code>start</code>	Starts the agent as a daemon process.
<code>stop</code>	Stops the agent's JVM process.
<code>restart</code>	Stops and then starts the agent's JVM process.
<code>status</code>	Queries the status of the agent's JVM process.
<code>dump</code>	Runs a thread dump for the agent process, and writes the result to the <code>agent.log</code> file in <code>AgentHome/log</code> .
<code>ping</code>	Pings the agent process.
<code>setup</code>	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
<code>install</code>	Installs the agent NT service. You must run <code>start</code> after running <code>install</code> .
<code>start</code>	Starts the agent as an NT service.
<code>stop</code>	Stops the agent as an NT service.

Option	Description
remove	Removes the agent's service from the NT service table.
query	Queries the current status of the agent NT service (status).
dump	Runs a thread dump for the agent process, and writes the result to the <code>agent.log</code> file in <code>AgentHome/log</code> .
ping	Pings the agent process.
setup	Re-registers the certificate using the existing token.

Managing an End Point Operations Management Agent on a Cloned Virtual Machine

When you clone a virtual machine that is running an End Point Operations Management agent that is collecting data, there are processes that you must complete related to data continuity to ensure data continuity.

Cloning a Virtual Machine to Delete the Original Virtual Machine

If you are cloning the virtual machine so that you can delete the original virtual machine, you need to verify that the original machine is deleted from the vCenter Server and from vRealize Operations Manager so that the new operating system to virtual machine relationship can be created.

Cloning a Virtual Machine to Run Independently of the Original Machine

If you are cloning the virtual machine so that you can run the two machines independently of the other, the cloned machine requires a new agent because an agent can only monitor a single machine.

Procedure

- ◆ On the cloned machine, delete the End Point Operations Management token and the `data` folder, according to the operating system of the machine.

Operating System	Process
Linux	Stop the End Point Operations Management services and delete the End Point Operations Management token and the <code>data</code> folder.
Windows	<ol style="list-style-type: none"> 1 Run <code>epops-agent remove</code>. 2 Remove the agent token and the <code>data</code> folder. 3 Run <code>epops-agent install</code>. 4 Run <code>epops-agent start</code>.

Moving Virtual Machines between vCenter Server Instances

When you move a virtual machine from one vCenter Server to another, vRealize Operations Manager preserves the unique object ID, identifiers, and historical data without creating any duplicate resources. This enables the new operating system to create a relationship with the migrated virtual machine.

Understanding Agent Uninstallation and Reinstallation Implications

When you uninstall or reinstall an End Point Operations Management agent, various elements are affected, including existing metrics that the agent has collected, and the identification token that enables a reinstalled agent to report on the previously discovered objects on the server. To ensure that you maintain data continuity, it is important that you are aware of the implications of uninstalling and reinstalling an agent.

There are two key locations related to the agent that are preserved when you uninstall an agent. Before reinstalling the agent, you must decide whether to retain or delete the files.

- The `/data` folder is created during agent installation. It contains the keystore, unless you chose a different location for it, and other data related to the currently installed agent.
- The `epops-token` platform token file is created before agent registration and is stored as follows:
 - Linux: `/etc/vmware/epops-token`
 - Windows: `%PROGRAMDATA%/VMware/EP Ops Agent/epops-token`

When you uninstall an agent, you must delete the `/data` folder. This does not affect data continuity.

However, to enable data continuity it is important that you do not delete the `epops-token` file. This file contains the identity token for the platform object. Following agent reinstallation, the token enables the agent to be synchronized with the previously discovered objects on the server.

When you reinstall the agent, the system notifies you whether it found an existing token, and provides its identifier. If a token is found, the system uses that token. If a token is not found, the system creates a new one. In the case of an error, the system prompts you to provide either a location and file name for the existing token file, or a location and file name for a new one.

The method that you use to uninstall an agent depends on how it was installed.

- [Uninstall an Agent that was Installed from an Archive](#)

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from an archive.
- [Uninstall an Agent that was Installed Using an RPM Package](#)

You can use this procedure to uninstall agents that you installed on virtual machines in your environment using an RPM package.
- [Uninstall an Agent that was Installed Using a Windows Executable](#)

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from a Windows `EXE` file.
- [Reinstall an Agent](#)

If you change the IP address, hostname or port number of the vRealize Operations Manager server, you need to uninstall and reinstall your agents.

Uninstall an Agent that was Installed from an Archive

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from an archive.

Prerequisites

Verify that the agent is stopped.

Procedure

- 1 (Optional) If you have a Windows operating system, run `ep-agent.bat remove` to remove the agent service.
- 2 Select the uninstall option that is appropriate to your situation.
 - If you do not intend to reinstall the agent after you have uninstalled it, delete the agent directory.
The default name of the directory is `epops-agent-version`.
 - If you are reinstalling the agent after you have uninstalled it, delete the `/data` directory.
- 3 (Optional) If you do not intend to reinstall the agent after you have uninstalled it, or you do not need to maintain data continuity, delete the `epops-token` platform token file.

Depending on your operating system, the file to delete is one of the following, unless otherwise defined in the properties file.

- Linux: `/etc/epops/epops-token`
- Windows: `%PROGRAMDATA%/VMware/EP Ops Agent/epops-token`

Uninstall an Agent that was Installed Using an RPM Package

You can use this procedure to uninstall agents that you installed on virtual machines in your environment using an RPM package.

When you are uninstalling an End Point Operations Management agent, it is good practice to stop the agent running, to reduce unnecessary load on the server.

Procedure

- ◆ On the virtual machine from which you are removing the agent, open a command line and run `rpm -e epops-agent`.

Results

The agent is uninstalled from the virtual machine.

Uninstall an Agent that was Installed Using a Windows Executable

You can use this procedure to uninstall agents that you installed on virtual machines in your environment from a Windows EXE file.

When you are uninstalling an End Point Operations Management agent, it is good practice to stop the agent running, to reduce unnecessary load on the server.

Procedure

- ◆ Double-click `unins000.exe` in the installation destination directory for the agent.

Results

The agent is uninstalled from the virtual machine.

Reinstall an Agent

If you change the IP address, hostname or port number of the vRealize Operations Manager server, you need to uninstall and reinstall your agents.

Prerequisites

To maintain data continuity, you must have retained the `epops-token` platform token file when you uninstalled your agent. See [Uninstall an Agent that was Installed from an Archive](#).

When you reinstall an End Point Operations Management agent on a virtual machine, objects that had previously been detected are no longer monitored. To avoid this situation, do not restart the End Point Operations Management agent until the plug-in synchronization is complete.

Procedure

- ◆ Run the agent install procedure that is relevant to your operating system.

See [Selecting an Agent Installer Package](#).

What to do next

After you reinstall an agent, MSSQL resources might stop receiving data. If this happens, edit the problematic resources and click **OK**.

Install Multiple End Point Operations Management Agents Simultaneously

If you have multiple End Point Operations Management agents to install at one time, you can create a single standardized `agent.properties` file that all the agents can use.

Installing multiple agents entails a number of steps. Perform the steps in the order listed.

Prerequisites

Verify that the following prerequisites are satisfied.

- 1 Set up an installation server.

An installation server is a server that can access the target platforms from which to perform remote installation.

The server must be configured with a user account that has permissions to SSH to each target platform without requiring a password.

- 2 Verify that each target platform on which an End Point Operations Management agent will be installed has the following items.
 - A user account that is identical to that created on the installation server.
 - An identically named installation directory, for example `/home/epomagent`.

- A trusted keystore, if required.

Procedure

1 [Create a Standard End Point Operations Management Agent Properties File](#)

You can create a single properties file that contains property values that multiple agents use.

2 [Deploy and Start Multiple Agents One-By-One](#)

You can perform remote installations to deploy multiple agents that use a single `agent.properties` file one-by-one.

3 [Deploy and Start Multiple Agents Simultaneously](#)

You can perform remote installations to simultaneously deploy agents that use a single `agent.properties` file.

Create a Standard End Point Operations Management Agent Properties File

You can create a single properties file that contains property values that multiple agents use.

To enable multiple agent deployment, you create an `agent.properties` file that defines the agent properties required for the agent to start up and connect with the vRealize Operations Manager server. If you supply the necessary information in the properties file, each agent locates its setup configuration at startup, rather than prompting you for the location. You can copy the agent properties file to the agent installation directory, or to a location available to the installed agent.

Prerequisites

Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.

Procedure

1 Create an `agent.properties` file in a directory.

You will copy this file later to other machines.

2 Configure the properties as required.

The minimum configuration is the IP address, user name, password, thumb print, and port of the vRealize Operations Manager installation server.

3 Save your configurations.

Results

The first time that the agents are started, they read the `agent.properties` file to identify the server connection information. The agents connect to the server and register themselves.

What to do next

Perform remote agent installations. See [Deploy and Start Multiple Agents One-By-One](#) or [Deploy and Start Multiple Agents Simultaneously](#).

Deploy and Start Multiple Agents One-By-One

You can perform remote installations to deploy multiple agents that use a single `agent.properties` file one-by-one.

Prerequisites

- Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.
- Verify that you configured a standard agent properties file and copied it to the agent installation, or to a location available to the agent installation.

Procedure

- 1 Log in to the installation server user account that you configured with permissions to use SSH to connect to each target platform without requiring a password.
- 2 Use SSH to connect to the remote platform.
- 3 Copy the agent archive to the agent host.
- 4 Unpack the agent archive.
- 5 Copy the `agent.properties` file to the `AgentHome/conf` directory of the unpacked agent archive on the remote platform.
- 6 Start the new agent.

Results

The agent registers with the vRealize Operations Manager server and the agent runs an autodiscovery scan to discover its host platform and supported managed products that are running on the platform.

Deploy and Start Multiple Agents Simultaneously

You can perform remote installations to simultaneously deploy agents that use a single `agent.properties` file.

Prerequisites

- Verify that the prerequisites in [Install Multiple End Point Operations Management Agents Simultaneously](#) are satisfied.
- Verify that you configured a standard agent properties file and copied it to the agent installation, or to a location available to the agent installation. See [Create a Standard End Point Operations Management Agent Properties File](#).

Procedure

- 1 Create a `hosts.txt` file on your installation server that maps the hostname to the IP address of each platform on which you are installing an agent.
- 2 Open a command-line shell on the installation server.

- 3 Type the following command in the shell, supplying the correct name for the agent package in the export command.

```
$ export AGENT=epops-agent-x86-64-linux-1.0.0.tar.gz
$ export PATH_TO_AGENT_INSTALL=</path/to/agent/install>
$ for host in `cat hosts.txt`; do scp $AGENT $host:$PATH_TO_AGENT_INSTALL && ssh $host "cd
$PATH_TO_AGENT_INSTALL; tar 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 discovers some of the objects to monitor. You can manually add other objects, such as files, scripts or processes, and specify the details so that the agent can monitor them.

The **Monitor OS Object** action only appears in the **Actions** menu of an object that can be a parent object.

Procedure

- 1 In the left pane of vRealize Operations Manager, select the agent adapter object that is to be the parent under which you are creating an OS object.

- 2 Select **Actions > Monitor OS Object**.

A list of parent object context-sensitive objects appear in the menu.

- 3 Choose one of the following options.

- Click an object type from the list to open the Monitor OS Object dialog box for that object type.

The three most popularly selected object types appear in the list.

- If the object type that you want to select is not in the list, click **More** to open the Monitor OS Object dialog box. Select the object type from the complete list of objects that are available for selection in the **Object Type** menu.

- 4 Specify a display name for the OS object.

- 5 Enter the appropriate values in the other text boxes.

The options in the menu are filtered according to the OS object type that you select.

Some text boxes might display default values, which you can overwrite if necessary. Note the following information about default values.

Option	Value
Process	<p>Supply the PTQL query in the form: <code>Class.Attribute.operator=value</code>.</p> <p>For example, <code>Pid.PidFile.eq=/var/run/sshd.pid</code>.</p> <p>Where:</p> <ul style="list-style-type: none"> ■ <code>Class</code> is the name of the Sigar class without the Proc prefix. ■ <code>Attribute</code> is an attribute of the given Class, index into an array or key in a Map class. ■ <code>operator</code> is one of the following (for String values): <ul style="list-style-type: none"> ■ <code>eq</code> Equal to value ■ <code>ne</code> Not Equal to value ■ <code>ew</code> Ends with value ■ <code>sw</code> Starts with value ■ <code>ct</code> Contains value (substring) ■ <code>re</code> Regular expression value matches <p>Delimit queries with a comma.</p>
Windows Service	<p>Monitor an application that runs as a service under Windows.</p> <p>To configure it, you supply its Service Name in Windows.</p> <p>To determine the Service Name:</p> <ol style="list-style-type: none"> 1 Select Run from the Windows Start menu. 2 Type <code>services.msc</code> in the run dialog box and click OK. 3 In the list of services displayed, right-click the service to monitor and choose Properties. 4 Locate the Service Name on the General tab.
Script	<p>Configure vRealize Operations Manager to periodically run a script that collects a system or application metric.</p>

6 Click **OK**.

You cannot click **OK** until you enter values for all the mandatory text boxes.

Results

The OS object appears under its parent object and monitoring begins.

Caution If you enter invalid details when you create an OS object, the object is created but the agent cannot discover it, and metrics are not collected.

Managing Objects with Missing Configuration Parameters

Sometimes when an object is discovered by vRealize Operations Manager for the first time, the absence of values for some mandatory configuration parameters is detected. You can edit the object's parameters to supply the missing values.

If you select **Custom Groups > Objects with Missing Configuration (EP Ops)** in the Environment Overview view of vRealize Operations Manager, you can see the list of all objects that have missing mandatory configuration parameters. In addition, objects with such missing parameters return an error in the Collection Status data.

If you select an object in the vRealize Operations Manager user interface that has missing configuration parameters, the red Missing Configuration State icon appears on the menu bar. When you point to the icon, details about the specific issue appear.

You can add the missing parameter values through the **Action > Edit Object** menu.

Mapping Virtual Machines to Operating Systems

You can map your virtual machines to an operating system to provide additional information to assist you to determine the root cause of why an alert was triggered for a virtual machine.

vRealize Operations Manager monitors your ESXi hosts and the virtual machines located on them. When you deploy an End Point Operations Management agent, it discovers the virtual machines and the objects that are running on them. By correlating the virtual machines discovered by the End Point Operations Management agent with the operating systems monitored by vRealize Operations Manager you have more details to determine the exact cause of an alert being triggered.

Verify that you have the vCenter Adapter configured with the vCenter Server that manages the virtual machines. You also need to ensure that you have VMware Tools that are compatible with the vCenter Server installed on each of the virtual machines.

User Scenario

vRealize Operations Manager is running but you have not yet deployed the End Point Operations Management agent in your environment. You configured vRealize Operations Manager to send you alerts when CPU problems occur. You see an alert on your dashboard because insufficient CPU capacity is available on one of your virtual machines that is running a Linux operating system. You deploy another two virtual CPUs but the alert remains. You struggle to determine what is causing the problem.

In the same situation, if you deployed the End Point Operations Management agent, you can see the objects on your virtual machines, and determine that an application-type object is using all available CPU capacity. When you add more CPU capacity, it also uses that. You disable the object and your CPU availability is no longer a problem.

Viewing Objects on Virtual Machines

After you deploy an End Point Operations Management agent on a virtual machine, the machine is mapped to the operating system and you can see the objects on that machine.

All the actions and the views that are available to other objects in your vRealize Operations Manager environment are also available for newly discovered server, service, and application objects, and for the deployed agent.

You can see the objects on a virtual machine in the inventory when you select the machine by clicking **Environment** from the menu, and then from the left pane click **vSphere Environment > vSphere Hosts and Clusters**. You can see the objects and the deployed agent under the operating system.

When you select an object, the center pane of the user interface displays data relevant to that objects.

Customizing How End Point Operations Management Monitors Operating Systems

End Point Operations Management gathers operating system metrics through agent-based collections. In addition to the features available after initial configuration of End Point Operations Management, you can enable remote monitoring, enable or disable plug-ins for additional monitoring, and customize End Point Operations Management logging.

Configuring Remote Monitoring

With remote monitoring you can monitor the state of an object from a remote location by configuring a remote check.

You can configure remote monitoring using HTTP, ICMP TCP methods.

When you configure a remote HTTP, ICMP or TCP check, it is created as a child object of the tested object that you are monitoring and of the monitoring agent.

If the object that you select to remotely monitor does not already have an alert configured, one is created automatically in the format `Remote check type failed on a object type`. If the object has an existing alert, that is used.

Configure Remote Monitoring of an Object

Use this procedure to configure remote monitoring of an object.

Configuration options are defined in [HTTP Configuration Options](#), [ICMP Configuration Options](#) and [TCP Configuration Options](#). You might need to refer to this information when you are completing this procedure.

Procedure

- 1 In the vRealize Operations Manager user interface, select the remote object to monitor.
- 2 On the details page for the object, select **Monitor this Object Remotely** from the **Actions** menu.
- 3 In the Monitor Remote Object dialog, select the End Point Operations Management agent that will remotely monitor the object from the **Monitored From** menu.
- 4 Select the method with which the remote object will be monitored from the **Check Method** menu.

The relevant parameters for the selected object type appear.

- 5 Enter values for all of the configuration options and click **OK**.

HTTP Configuration Options

Here are the options in the configuration schema for the HTTP resource.

For the HTTP resource, the `netsservices` plug-in descriptor default values are:

- `port: 80`

■ `sslport: 443`

HTTP Configuration Options

Table 4-68. ssl Option

Option Information	Value
Description	Use ssl
Default	false
Optional	true
Type	boolean
Notes	N/A
Parent Schema	ssl

Table 4-69. 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 4-70. port Option

Option Information	Value
Description	Port
Default	A default value for port is set for each type of network service by properties in the netsservices plug-in descriptor.
Optional	false
Type	N/A
Notes	The port on which the service listens.
Parent Schema	sockaddr

Table 4-71. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	true
Type	int
Notes	The maximum length of time the agent waits for a response to a request to the remote service.
Parent Schema	sockaddr

Table 4-72. 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 4-73. 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 4-74. hostheader Option

Option Information	Value
Description	Host Header
Default	none
Optional	true
Type	N/A
Notes	Use this option to set a <code>Host</code> 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, <code>blog.mypost.com</code> .
Parent Schema	http

Table 4-75. follow Option

Option Information	Value
Description	Follow Redirects
Default	enabled
Optional	true
Type	boolean
Notes	Enable if the HTTP request that is generated will be redirected. This is important, because an HTTP server returns a different code for a redirect and vRealize Operations Manager determines that the HTTP service check is unavailable if it is a redirect, unless this redirect configuration is set.
Parent Schema	http

Table 4-76. pattern Option

Option Information	Value
Description	Response Match (substring or regex)
Default	none
Optional	true
Type	N/A

Table 4-76. pattern Option (continued)

Option Information	Value
Notes	Specify a pattern or substring for vRealize Operations Manager to attempt to match against the content in the HTTP response. This enables you to check that in addition to being available, the resource is serving the content you expect.
Parent Schema	http

Table 4-77. 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 4-78. 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 4-79. Credential Option

Option Information	Value
Description	Username
Default	N/A
Optional	true
Type	N/A

Table 4-79. Credential Option (continued)

Option Information	Value
Notes	Supply the user name if the target site is password-protected.
Parent Schema	credentials

ICMP Configuration Options

Here are the options in the configuration schema for the ICMP resource.

ICMP configuration is not supported in Windows environments. When attempting to run an ICMP check for remote monitoring from an Agent running on a Windows platform, no data is returned.

Table 4-80. 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	netsservices plug-in descriptor

Table 4-81. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A
Type	int
Notes	The maximum time period the agent waits for a response to a request to the remote service.
Parent Schema	netsservices plug-in descriptor

TCP Configuration Options

Here are the options in the configuration schema to enable TCP checking.

Table 4-82. port Option

Option Information	Value
Description	Port
Default	A default value for port is set for each type of network service by properties in the netservices plug-in descriptor.
Optional	false
Type	N/A
Notes	The port on which the service listens.
Parent Schema	sockaddr

Table 4-83. hostname Option

Option Information	Value
Description	Hostname
Default	localhost
Optional	N/A
Type	N/A
Notes	The hostname of system that hosts the object to monitor. For example: mysite.com
Parent Schema	netservices plug-in descriptor

Make sure that you use the IP address of the machine on which the remote check is to run, not the host name.

Table 4-84. sotimeout Option

Option Information	Value
Description	Socket Timeout (in seconds)
Default	10
Optional	N/A
Type	int
Notes	The maximum amount of time the agent waits for a response to a request to the remote service.
Parent Schema	netservices plug-in descriptor

Agent Management

You can add, edit, and delete End Point Operations Management agents and enable or disable the End Point Operations Management plug-ins from the tabs in the Agent Management page.

Where You Find the Agent Management Page

In the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**.

Agents Tab

You can view the End Point Operations Management agents that are installed and deployed in your environment.

Where You Find the Agents Tab

In the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**.

How the Agents Tab Works

You can view all the agents that are installed, the virtual machines on which they are installed, their operating system and the agent bundle version. You can also view the collection details of each agent. You can filter the list of agents based on the name of the agent. You add a filter from the upper-right corner of the toolbar. You can sort the Agent Token, Agent Name, Collection State, and Collection Status columns by clicking the column name.

Plug-ins Tab

End Point Operations Management agents include plug-ins that determine which objects to monitor, how they should be monitored, which metrics to collect, and so on. Some plug-ins are included in the default End Point Operations Management agent installation, and other plug-ins might be added as part of any management pack solution that you install to extend the vRealize Operations Manager monitoring process.

You can use the **Plug-ins** tab from the Agents Management page to disable or enable the agent plug-ins that are deployed in your environment as part of a solution installation. For example, you might want to temporarily disable a plug-in so that you can analyze the implication of that plug-in on a monitored virtual machine. To access the **Plug-ins** tab, in the menu, click **Administration**, and then in the left pane click **Configuration > End Point Operations**. You can sort all the columns in the tab by clicking the column name.

All the default plug-ins and the plug-ins that are deployed when you installed one or more solutions are listed alphabetically on the tab.

You must have Manage Plug-ins permissions to enable and disable plug-ins.

When you disable a plug-in, it is removed from all the agents on which it has existed, and the agent no longer collects the metrics and other data related to that plug-in. The plug-in is marked as disabled on the vRealize Operations Manager server.

You cannot disable the default plug-ins that are installed during the vRealize Operations Manager installation.

You use the action menu that appears when you click the gear wheel icon to disable or enable plug-ins.

Before you deploy a new version of a plug-in, you must implement a shutdown method. If you do not implement a shutdown method, the existing plug-in version does not shut down so that a new instance is created and allocated resources such as static threads are not released. Implement a shutdown method for these plug-ins.

- Plug-ins that use third-party libraries
- Plug-ins that use native libraries
- Plug-ins that use connection pools
- Plug-ins that might lock files, which cause issues on Windows operating systems

It is good practice that plug-ins do not use threads, third-party libraries, or static collection.

Configuring Plug-in Loading

At startup, an End Point Operations Management agent loads all the plug-ins in the `AgentHome/bundles/agent-x.y.z-nnnn/pdk/plugins` directory. You can configure properties in the `agent.properties` file to reduce an agent's memory footprint by configuring it to load only the plug-ins that you use.

Plug-ins are deployed to all agents when a solution is installed. You might want to use the properties described here in a situation in which you need to remove one or more plug-ins from a specific machine. You can either specify a list of plug-ins to exclude, or configure a list of plug-ins to load.

`plugins.exclude`

Use this property to specify the plug-ins that the End Point Operations Management agent must not load at startup.

You supply a comma-separated list of plugins to exclude. For example,

```
plugins.exclude=jboss,apache,mysql.
```

`plugins.include`

Use this property to specify the plug-ins that the End Point Operations Management agent must load at startup.

You supply a comma-separated list of plugins to include. For example,

```
plugins.include=weblogic,apache.
```

Understanding the Unsynchronized Agents Group

An unsynchronized agent is an agent that is not synchronized with the vRealize Operations Manager server in terms of its plug-ins. The agent might be missing plug-ins that are registered on the server, include plug-ins that are not registered on the server, or include plug-ins that have a different version to that registered on the server.

Each agent must be synchronized with the vRealize Operations Manager server. During the time that an agent is not synchronized with the server, it appears in the Unsynchronized Agents list. The list is located in the vRealize Operations Manager user interface on the **Groups** tab in the Environment view.

The first time an agent is started, a status message is sent to the server. The server compares the status sent by the agent with that on the server. The server sends commands to the agent to synchronize, download or delete plug-ins, as required by the differences that it detects.

When a plug-in is deployed, disabled, or enabled as part of a management pack solution update, the vRealize Operations Manager server detects that change and sends a new command to the agents so that synchronization occurs.

Commonly, multiple agents are affected at the same time when a plug-in is deployed, disabled or enabled. All agents have an equal need to be updated so, to avoid overloading the server and creating performance issues that might occur if many agents were all synchronized at the same time, synchronization is performed in batches and is staggered in one-minute periods. You will notice that the list of unsynchronized agents decrements over time.

Configuring Agent Logging

You can configure the name, location, and logging level for End Point Operations Management agent logs. You can also redirect system messages to the agent log, and configure the debug log level for an agent subsystem.

Agent Log Files

The End Point Operations Management agent log files are stored in the `AgentHome/log` directory.

Agent log files include the following:

`agent.log`

`agent.operations.log`

This log is applicable to Windows-based agents only.

This is an audit log that records the commands that were run on the agent, together with the parameters that the agent used to action them.

`wrapper.log`

The Java service wrapper-based agent launcher writes messages to the `wrapper.log` file. For a non-JRE agent, this file is located in `agentHome/wrapper/sbin`.

In the event that the value was changed ifr the `agent.logDir` property, the file is also located in `agentHome/wrapper/sbin`.

Configuring the Agent Log Name or Location

Use these properties to change the name or location of the agent log file.

`agent.logDir`

You can add this property to the `agent.properties` file to specify the directory where the End Point Operations Management agent will write its log file. If you do not specify a fully qualified path, `agent.logDir` is evaluated relative to the agent installation directory.

This property does not exist in the `agent.properties` file unless you explicitly add it. The default behavior is equivalent to the `agent.logDir=log` setting, resulting in the agent log file being written to the `AgentHome/log` directory.

To change the location for the agent log file, add `agent.logDir` to the `agent.properties` file and enter a path relative to the agent installation directory, or a fully qualified path.

The name of the agent log file is configured with the `agent.logFile` property.

`agent.logFile`

This property specifies the path and name of the agent log file.

In the `agent.properties` file, the default setting for the `agent.LogFile` property is made up of a variable and a string, `agent.logFile=${agent.logDir}\agent.logDir`.

- *agent.logDir* is a variable that supplies the value of an identically named agent property. By default, the value of *agent.logDir* is `log`, interpreted relative to the agent installation directory.
- `agent.log` is the name for the agent log file.

By default, the agent log file is named `agent.log` and is written to the `AgentHome/log` directory.

To configure the agent to log to a different directory, you must explicitly add the `agent.logDir` property to the `agent.properties` file.

Configuring the Agent Logging Level

Use this property to control the severity level of messages that the End Point Operations Management agent writes to the agent log file.

`agent.logLevel`

This property specifies the level of detail of the messages that the End Point Operations Management agent writes to the log file.

Setting the `agent.logLevel` property value to `DEBUG` level is not advised. This level of logging across all subsystems imposes overhead, and can also cause the log file to roll over so frequently that log messages of interest are lost. It is preferable to configure debug level logging only at the subsystem level.

The changes that you make to this property become effective approximately five minutes after you save the properties file. It is not necessary to restart the agent to initiate the change.

Redirecting System Messages to the Agent Log

You can use these properties to redirect system-generated messages to the End Point Operations Management agent log file.

`agent.logLevel.SystemErr`

This property redirects `System.err` to `agent.log`. Commenting out this setting causes `System.err` to be directed to `agent.log.startup`.

The default value is `ERROR`.

`agent.logLevel.SystemOut`

This property redirects `System.out` to `agent.log`. Commenting out this setting causes `System.out` to be directed to `agent.log.startup`.

The default value is `INFO`.

Configuring the Debug Level for an Agent Subsystem

For troubleshooting purposes, you can increase the logging level for an individual agent subsystem.

To increase the logging level for an individual agent subsystem, uncomment the appropriate line in the section of the `agent.properties` file that is labeled `Agent Subsystems:` Uncomment individual subsystems to see debug messages.

Agent log4j Properties

This is the `log4j` properties in the `agent.properties` file.

```
log4j.rootLogger=${agent.logLevel}, R

log4j.appender.R.File=${agent.logFile}
log4j.appender.R.MaxBackupIndex=1
log4j.appender.R.MaxFileSize=5000KB
log4j.appender.R.layout.ConversionPattern=%d{dd-MM-yyyy HH:mm:ss,SSS z} %-5p [%t] [%c{1}]@%L]
%m%n
log4j.appender.R.layout=org.apache.log4j.PatternLayout
log4j.appender.R=org.apache.log4j.RollingFileAppender

##
## Disable overly verbose logging
##
log4j.logger.org.apache.http=ERROR
log4j.logger.org.springframework.web.client.RestTemplate=ERROR
log4j.logger.org.hyperic.hq.measurement.agent.server.SenderThread=INFO
log4j.logger.org.hyperic.hq.agent.server.AgentDListProvider=INFO
log4j.logger.org.hyperic.hq.agent.server.MeasurementSchedule=INFO
log4j.logger.org.hyperic.util.units=INFO
log4j.logger.org.hyperic.hq.product.pluginxml=INFO

# Only log errors from naming context
log4j.category.org.jnp.interfaces.NamingContext=ERROR
log4j.category.org.apache.axis=ERROR

#Agent Subsystems: Uncomment individual subsystems to see debug messages.
#-----
#log4j.logger.org.hyperic.hq.autoinventory=DEBUG
#log4j.logger.org.hyperic.hq.livedata=DEBUG
#log4j.logger.org.hyperic.hq.measurement=DEBUG
#log4j.logger.org.hyperic.hq.control=DEBUG

#Agent Plugin Implementations
#log4j.logger.org.hyperic.hq.product=DEBUG

#Server Communication
#log4j.logger.org.hyperic.hq.bizapp.client.AgentCallbackClient=DEBUG

#Server Realtime commands dispatcher
#log4j.logger.org.hyperic.hq.agent.server.CommandDispatcher=DEBUG

#Agent Configuration parser
#log4j.logger.org.hyperic.hq.agent.AgentConfig=DEBUG
```



```
#Agent plugins loader
#log4j.logger.org.hyperic.util.PluginLoader=DEBUG

#Agent Metrics Scheduler (Scheduling tasks definitions & executions)
#log4j.logger.org.hyperic.hq.agent.server.session.AgentSynchronizer.SchedulerThread=DEBUG

#Agent Plugin Managers
#log4j.logger.org.hyperic.hq.product.MeasurementPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.AutoinventoryPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ConfigTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LogTrackPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.LiveDataPluginManager=DEBUG
#log4j.logger.org.hyperic.hq.product.ControlPluginManager=DEBUG
```

NSX-T

The NSX-T adapter allows you to retrieve alerts and findings from NSX-T to vRealize Operations Manager.

The NSX-T adapter supports adapter configuration using vIDM for NSX-T versions 3.0 and above. The roles and permissions associated with the vIDM users collecting the NSX-T adapter data is:

Roles	Permissions
Enterprise Admin	Collect all data.
VPN admin	Collect only Management appliance and NSX cluster data.
Network engineer	<ul style="list-style-type: none"> Collect all the NSX-T resources except the Load Balancer and collect limited routers data. Router data collected: <ul style="list-style-type: none"> Tier 0 router connected to logical switch. Tier 1 router created from vCloud Director.
<ul style="list-style-type: none"> Security Engineer Security Operator Auditor 	Collect all data except the load balancer.
<ul style="list-style-type: none"> LB Admin LB Auditor Netxpartner Admin 	Cannot collect any data.

Configuring the NSX-T Adapter

Procedure

- 1 On the menu, click **Administration** and in the left pane, click **Solutions > Other Accounts**.
- 2 On the Other Accounts page, click **Add Account**.
- 3 On the Account types page, click **NSX-T Adapter**.

- 4 Enter a display name and description for the NSX-T account.
 - Name. Enter the name for the NSX-T instance as you want it to appear in vRealize Operations Manager.
 - Description. Enter any additional information that helps you manage your instances.
- 5 Virtual IP/NSX-T Manager. Enter the FQDN, the IP address, or the Virtual IP of the NSX-T manager.
- 6 Select the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the NSX-T environment , click the plus sign.
 - Credential Name. The name by which you are identifying the configured credentials.
 - User Name. The user name of the NSX-T instance.
 - Password. The password of the NSX-T instance.
- 7 Determine which vRealize Operations Manager collector or collector group is used to manage the account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.
- 8 Click **Validate Connection** to validate the connection.
- 9 To save the configurations, click **Save This SDDC**.
- 10 Click **Add**.

The adapter instance is added to the list.

What to do next

Verify that the adapter is configured and collecting data.

Configuring Alerts and Actions

In vRealize Operations Manager , alerts and actions play key roles in monitoring the objects.

Triggered Alerts

The **Triggered Alerts** page is a list of all the alerts generated in vRealize Operations Manager . Use the alert list to determine the state of your environment and to begin resolving problems.

How the Triggered Alerts Page Works

By default, only active alerts are initially listed, and the alerts are grouped by Time. Review and manage the alerts in the list using the toolbar options. Select multiple rows in the list using Shift+click, Control+click.

To see the alert details, click the alert name. The alert details appear on the right, including the symptoms triggered by the alert. The system offers recommendations for addressing the alert and link to run the recommendation. A Run Action button may appear in the details. Hover over the button to learn what recommendation is performed if you click the button. Alternatively, you can view the **Run** button and the **Suggested Fix** in the Alerts data grid. You can filter by alerts that have the Run option enabled and perform the recommended task to address the alert from the Alerts data grid. Click the small box on the lower left of the alert list to include the **Suggested Fix** and **Run** columns in the data grid.

Click the name of the object on which the alert was generated to see the object details, and access additional information relating to metrics and events.

If you migrated alerts from a previous version of vRealize Operations Manager , the alerts are listed with a cancelled status and alert details are not available.

Where You Find the All Alerts Page

In the menu, click **Alerts**.

Triggered Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to sort the alert list and to cancel, suspend, or manage ownership. Use the data grid to view the alerts and alert details.

Select an alert from the list to enable the Actions menu:

Table 4-85. Actions Menu

Option	Description
Cancel Alert	<p>Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.</p> <p>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.</p>
Delete Canceled Alerts	<p>Delete cancelled (inactive) alerts by doing a group selection or by individually selecting alerts. The option is disabled for active alerts.</p>

Table 4-85. Actions Menu (continued)

Option	Description
Suspend	<p>Suspend an alert for a specified number of minutes.</p> <p>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.</p> <p>The user who suspends the alert becomes the assigned owner.</p>
Assign to	Assign the alert to a user. You can search for a specific username and click Save to assign the alert to the selected user.
Take Ownership	<p>As the current user, you make yourself the owner of the alert.</p> <p>You can only take ownership of an alert, you cannot assign ownership.</p>
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...	<p>Offers two options for disabling the alert:</p> <p>Disable the alert in all policies: this disables the alert for all objects for all the policies.</p> <p>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.</p>
Open an external application	<p>Actions you can run on the selected object.</p> <p>For example, Open Virtual Machine in vSphere Client.</p>

Table 4-86. Group By Options

Option	Description
None	Alerts are not sorted into specific groupings.
Time	Group alerts by time triggered. This is the default option. You can also group by 1 hour, 4 hours, Today and Yesterday, days of current week, Last week and Older.
Criticality	Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the "All Alerts Data Grid Options" table, below.
Definition	Group alerts by definition, that is, group like alerts together.
Object Type	Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together.

Table 4-87. All Filters

All Filters	Descriptions
Filtering options	<p>Limit the list of alerts to those matching the filters you choose.</p> <p>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.</p>
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	<p>State of user interaction with the alert. Possible values include:</p> <ul style="list-style-type: none"> ■ Open. The alert is available for action and has not been assigned to a user. ■ Assigned. The alert is assigned to the user who is logged in when that user clicks Take Ownership. ■ Suspended. The alert was suspended for a specified amount of time. The alert is temporarily excluded from affecting the health, risk, and efficiency of the object. This state is useful when a system administrator is working on a problem and does not want the alert to affect the health status of the object.
Object Type	Type of object on which the alert was generated.
Updated On	<p>Date and time when the alert was last modified.</p> <p>An alert is updated whenever one of the following changes occurs:</p> <ul style="list-style-type: none"> ■ Another symptom in the alert definition is triggered. ■ Triggering symptom that contributed to the alert is canceled.

Table 4-87. All Filters (continued)

All Filters	Descriptions
Canceled On	<p>Date and time when the alert canceled for one of the following reasons:</p> <ul style="list-style-type: none"> ■ Symptoms that triggered the alert are no longer active. Alert is canceled by the system. ■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions are disabled in the policy that is applied to the object. ■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions were deleted. ■ Alert definition for this alert is disabled in the policy that is applied to the object. ■ Alert definition is deleted. ■ User canceled the alert.
Action	<p>Choose Yes to filter based on alerts that have the Run option enabled. Choose No to filter based on alerts that have the Run option disabled.</p>

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 4-88. Triggered Alerts Data Grid

Option	Description
Criticality	<p>Criticality is the level of importance of the alert in your environment.</p> <p>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.</p> <p>The possible values include:</p> <ul style="list-style-type: none"> ■ Critical ■ Immediate ■ Warning ■ Information
Alert	<p>Name of the alert definition that generated the alert.</p> <p>Click the alert name to display the alert details to the right.</p>
Triggered On	<p>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.</p> <p>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object.</p>
Created On	Date and time when the alert was generated.
Status	<p>Current state of the alert.</p> <p>Possible values include Active or Canceled.</p>

Table 4-88. Triggered 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.
Importance	Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm.
Suggested Fix	Displays the recommendation to address the alert.
Action	Click this button to perform the recommendation to address the alert.

Types of Alerts

Alerts in vRealize Operations Manager are of three types. The alert type determines the severity of the problem.

Health Alerts

The health alert list is all the generated alerts that are configured to affect the health of your environment and require immediate attention. You use the health alert list to evaluate, prioritize, and immediately begin resolving the problems.

Risk Alerts

The risk alerts list is all the generated alerts that are configured to indicate risk in your environment. Address risk alerts in the near future, before the triggering symptoms that generated the alert negatively affect the health of your environment.

Efficiency Alerts

The efficiency alerts list is all the generated alerts that are configured to indicate problems with the efficient use of your monitored objects in your environment. Address efficiency alerts to reclaim wasted space or to improve the performance of objects in your environment.

Alert Information

When you click an alert from the all alerts list, the alert information appears on the right. View the alert information to see the symptoms which triggered the alert, recommendations to fix the underlying issue, and troubleshoot the cause of the alert.

How You View the Alert Information

- In the menu, click **Alerts**. Click an alert from the alert list.
- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. Click the object and then the **Alerts** tab.
- In the menu, select Search and locate the object of interest. Click the object and then the **Alerts** tab.

The alert description is hidden when you open the alert information. Click **View Description** to see the description of the alert. View the time stamp of when the alert started, and when it was updated, below the alert title.

Alert Details Tab

Section	Description
Recommendations	View recommendations for the alert. Click < or > to cycle through the recommendations. To resolve the alert, click the Run Action button if it appears.
Other Recommendations	Collapse the section to view additional recommendations. See the links in the Need More Information? section to view additional metrics, events, or other details that appear as a link.
Symptoms	View the symptoms that triggered the alert. Collapse each symptom to view additional information.
Notes	Enter your notes about the alert and click Submit to save.
Close	Click the X icon to close the alert details tab.

Related Alerts Tab

The **Related Scope** displayed on the right, shows the objects that are one level above and one level below the object on which the alert was triggered. This topology is fixed. You cannot change the scope in the **Related Alerts** tab.

On the right, you can see the following:

- If the same alert was triggered on the object in the past 30 days. This helps you understand if this is a recurring problem or something new.
- If the same alert was triggered on other peers in the same environment, in the past 30 days. This helps you do a quick peer analysis to understand if others are impacted with the same problem.
- All the alerts triggered in the current topology. This helps you investigate if there are other alerts upstream or downstream in the environment which are impacting the health of the object.

Potential Evidence Tab

See the **Potential Evidence** tab for potential evidences around the problem, and to arrive at the root cause. This tab displays events, property changes, and anomalous metrics potentially relevant to the alert. The time range and the scope are fixed. To modify the scope or the time range and investigate further, click **Launch Workbench**. This runs the troubleshooting workbench.

The time range that is displayed in the potential evidence tab is two hours and thirty minutes before the alert was triggered. vRealize Operations Manager looks for potential evidences in this time range.

Configuring Alerts

Whenever there is a problem in the environment, the alerts are generated. You can create the alert definitions so that the generated alerts tell you about the problems in the monitored environment.

Defining Alerts in vRealize Operations Manager

An alert definition comprises one or more symptom definitions, and the alert definition is associated with a set of recommendations and actions that help you resolve the problem. Alert definitions include triggering symptom definitions and actionable recommendations. You create the alert definitions so that the generated alerts tell you about problems in the monitored environment. You can then respond to the alerts with effective solutions that are provided in the recommendations.

Predefined alerts are provided in vRealize Operations Manager as part of your configured adapters. You can add or modify alert definitions to reflect the needs of your environment.

Symptoms in Alert Definitions

Symptom definitions evaluate conditions in your environment that, if the conditions become true, trigger a symptom and can result in a generated alert. You can add symptom definitions that are based on metrics or super metrics, properties, message events, fault events, or metric events. You can create a symptom definition as you create an alert definition or as an individual item in the appropriate symptom definition list.

When you add a symptom definition to an alert definition, it becomes a part of a symptom set. A symptom set is the combination of the defined symptom with the argument that determines when the symptom condition becomes true.

A symptom set combines one or more symptom definitions by applying an Any or All condition, and allows you to choose the presence or absence of a particular symptom. If the symptom set pertains to related objects rather than to Self, you can apply a population clause to identify a percentage or a specific count of related objects that exhibit the included symptom definitions.

An alert definition comprises one or more symptom sets. If an alert definition requires all of the symptom sets to be triggered before generating an alert, and only one symptom set is triggered, an alert is not generated. If the alert definition requires only one of several symptom sets to be triggered, then the alert is generated even though the other symptom sets were not triggered.

Recommendations in Alert Definitions

Recommendations are the remediation options that you provide to your users to resolve the problems that the generated alert indicates.

When you add an alert definition that indicates a problem with objects in your monitored environment, add a relevant recommendation. Recommendations can be instructions to your users, links to other information or instruction sources, or vRealize Operations Manager actions that run on the target systems.

Modifying Alert Definitions

If you modify the alert impact type of an alert definition, any alerts that are already generated will have the previous impact level. Any new alerts will be at the new impact level. If you want to reset all the generated alerts to the new level, cancel the old alerts. If they are generated after cancellation, they will have the new impact level.

Defining Symptoms for Alerts

Symptoms are conditions that indicate problems in your environment. You define symptoms that you add to alert definitions so that you know when a problem occurs with your monitored objects.

As data is collected from your monitored objects, the data is compared to the defined symptom condition. If the condition is true, then the symptom is triggered.

You can define symptoms based on metrics and super metrics, properties, message events, fault events, and metric events.

Defined symptoms in your environment are managed in the Symptom Definitions. When the symptoms that are added to an alert definition are triggered, they contribute to a generated alert.

Define Symptoms to Cover All Possible Severities and Conditions

Use a series of symptoms to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of `Warning` while `Volume reached capacity limit` might have a severity level of `Critical`. The first symptom is not an immediate threat. The second symptom is an immediate threat.

About Metrics and Super Metrics Symptoms

Metric and super metric symptoms are based on the operational or performance values that vRealize Operations Manager collects from target objects in your environment. You can configure the symptoms to evaluate static thresholds or dynamic thresholds.

You define symptoms based on metrics so that you can create alert definitions that let you know when the performance of an object in your environment is adversely affected.

Static Thresholds

Metric symptoms that are based on a static threshold compare the currently collected metric value against the fixed value you configure in the symptom definition.

For example, you can configure a static metric symptom where, when the virtual machine CPU workload is greater than 90, a critical symptom is triggered.

Dynamic Thresholds

Metric symptoms that are based on dynamic thresholds compare the currently collected metric value against the trend identified by vRealize Operations Manager , evaluating whether the current value is above, below, or generally outside the trend.

For example, you can configure a dynamic metric symptom where, when the virtual machine CPU workload is above the trended normal value, a critical symptom is triggered.

Metric / Super Metric Symptom Definitions

The Metric / Super Metric Symptom Definitions is a list of the metric-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined metric threshold triggering states and determine if you want to add, edit, or clone symptoms.

Where You Find Metric / Super Metric Symptoms

To manage symptoms based on metrics and super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-89. Metric / Super Metric Symptoms Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.</p> <ul style="list-style-type: none"> ■ Add. Add a symptom definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ 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	<p>Limits the list to symptoms matching the filter.</p> <p>You can also sort on the columns in the data grid.</p>

Table 4-89. Metric / Super Metric Symptoms Options (continued)

Option	Description
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.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Metric and Supermetric Symptoms Definition Workspace

You define metric and super metric symptoms, which are based on collected operational or performance values, so that you can create one or more of the symptoms that you can add to an alert definition in vRealize Operations Manager . When a symptom is triggered, you use the symptoms to evaluate alerts or troubleshoot other problems.

How Metric Symptom Definitions Work

A metric or super metric symptom is triggered when a metric is compared to the configured static or dynamic thresholds, and the symptom condition is evaluated as true. If the symptom is based on a static threshold, the metric is compared based on the configured operator and the provided numeric value. If the symptom is based on a dynamic threshold, the metric is compared based on whether the current value is above, below, or abnormal compared to the calculated trend value.

Where You Find the Metric Symptom Definition Workspace

To define symptoms based on metrics or super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric / Property**. Click **Add** to define a metric-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-90. Symptoms Workspace Options for Metrics and Super Metrics

Option	Description
Metric Explorer	Components that you use to locate your metrics or super metrics for which you are creating symptoms.
Base Object Type	Object against which the symptom is evaluated. Based on the select object type, the list of available metrics displays only the metrics applicable to the object type.
Select Resource	If a metric or supermetric is not listed in the common metric or supermetric list, based on the selected based object type, use Select Resource to inspect the metrics or supermetrics of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a metric or supermetric for a specific object, the symptom definition is applicable to all objects with that metric or supermetric in your environment.
Search	Use a word search to limit the number of items that appear in the list.
Metric list	List of metrics for the selected base object type.
Symptom definition workspace	Click and drag the metric to the right pane. You can define symptoms based on static or dynamic thresholds.
Threshold	<p>Determines if the symptom is static or dynamic.</p> <ul style="list-style-type: none"> ■ 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 4-90. Symptoms Workspace Options for Metrics and Super Metrics (continued)

Option	Description
Static Threshold configuration options	<p>If you select Static Threshold, configure the options for this threshold type.</p> <ul style="list-style-type: none"> ■ Operator. Determines how the value you specify in the value text box is compared to the current value of the metric or super metric when the symptom is evaluated. ■ Value. Value that is the triggering threshold. ■ Criticality level. Severity of the symptom when it is triggered. ■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms. ■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true. ■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false. ■ Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered. ■ Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking Select Object next to the Metrics text box.
Dynamic Threshold configuration options	<p>If you select Dynamic Threshold, configure the options for this threshold type.</p> <ul style="list-style-type: none"> ■ Threshold trend. Relationship of the current value to trended range based on the following options: <ul style="list-style-type: none"> ■ 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.

Table 4-90. Symptoms Workspace Options for Metrics and Super Metrics (continued)

Option	Description
	<ul style="list-style-type: none"> ■ 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 viewing triggered symptoms. ■ Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered. ■ Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking Select Object next to the Metrics field.

Property Symptoms

Property symptoms are based on the configuration properties that vRealize Operations Manager collects from the target objects in your environment.

You define symptoms based on properties so that you can create alert definitions that let you know when changes to properties on your monitored objects can affect the behavior of the objects in your environment.

Property Symptoms Definitions

The Property Symptom Definitions is a list of the property-based symptoms in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined property triggering states and determine whether to add, edit, or clone symptoms.

Where You Find Property Symptoms

To manage symptoms based on properties, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions > Metric/Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-91. Property Symptoms Definitions Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.</p> <ul style="list-style-type: none"> ■ Add. Add a symptom definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
All Filters	Limits the list to symptoms matching the filter. You can also sort on the columns in the data grid.
Quick Filter (Name)	Limits the list based on the text you type.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Property	Text string that is used as a reference key for the property. You can use the property to locate additional information about the property.
Operator	Operator used to compare the threshold value to the current value.
Value	Text string that is the compared value for the property.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Property Symptoms Definition Workspace

You define property symptoms, which are based on collected configuration properties, so that you can add one or more symptoms to an alert definition in vRealize Operations Manager . You use the triggered symptoms to resolve alerts or troubleshoot other problems.

How Property Symptom Definitions Work

A property symptom is triggered when the defined threshold is compared with the current property value and the comparison is evaluated as true.

Where You Find the Property Symptom Definition Workspace

To define symptoms based on metrics or super metrics, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-92. 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 4-92. Symptoms Workspace Options for Properties (continued)

Option	Description
Symptom definition workspace	Drag the property to the right pane.
Property	<p>The properties are configured values that are compared to the value you specify. You can configure a single property symptom or add multiple symptoms.</p> <p>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.</p> <p>Configure the options:</p> <ul style="list-style-type: none"> ■ Operator. Determines how the value you specify in the value text box is compared to the current value of the property for an object when the symptom definition is evaluated. ■ Value. Value that the operator evaluates. ■ Criticality level. Severity of the symptom when it is triggered. ■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms. ■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true. ■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

Message Event Symptoms

Message event symptoms are based on events received as messages from a component of vRealize Operations Manager or from an external monitored system through the system's REST API. You define symptoms based on message events to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

The adapters for the external monitored systems and the REST API are inbound channels for collecting events from external sources. Adapters and the REST server both run in the vRealize Operations Manager system. The external system sends the messages, and vRealize Operations Manager collects them.

You can create message event symptoms for the supported event types. The following list is of supported event types with example events.

- **System Performance Degradation.** This message event type corresponds to the `EVENT_CLASS_SYSTEM` and `EVENT_SUBCLASS_PERFORM_DEGRADATION` type and subtype in the vRealize Operations Manager API SDK.
- **Change.** The VMware adapter sends a change event when the CPU limit for a virtual machine is changed from unlimited to 2 GHz. You can create a symptom to detect CPU contention issues as a result of this configuration change. This message event type corresponds to the `EVENT_CLASS_CHANGE` and `EVENT_SUBCLASS_CHANGE` type and subtype in the vRealize Operations Manager API SDK.
- **Environment Down.** The vRealize Operations Manager adapter sends an environment down event when the collector component is not communicating with the other components. You can create a symptom that is used for internal health monitoring. This message event type corresponds to the `EVENT_CLASS_ENVIRONMENT` and `EVENT_SUBCLASS_DOWN` type and subtype in the vRealize Operations Manager API SDK.
- **Notification.** This message event type corresponds to the `EVENT_CLASS_NOTIFICATION` and `EVENT_SUBCLASS_EXTEVENT` type and subtype in the vRealize Operations Manager API SDK.

Message Event Symptom Definitions

The Message Event Symptom Definitions is a list of the message event-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined message events and to determine if you want to add, edit, or clone symptoms.

Where You Find Message Event Symptoms

To manage symptoms based on message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Message Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-93. Message Event Symptoms Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.</p> <ul style="list-style-type: none"> ■ Add. Add a symptom definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to symptoms matching the filter.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Event Type	Defined event classification type.
Operator	Operator used to compare the message from the incoming event against the event message specified in the symptom.
Event Message	Text string that is compared to the message in the incoming event using the specified operator.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Message Event Symptoms Definition Workspace

Message event symptoms are based on message events received from a component of vRealize Operations Manager or an external monitored system through the system's REST API. You define message event systems so that you can create one or more of the symptoms that you can add to an alert definition.

How Message Event Symptom Definitions Work

A message event symptom is triggered when a message in an incoming event matches the text string in the symptom, based on the specified operator.

Where You Find the Message Event Symptom Definition Workspace

To define symptoms based on message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Message Event** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-94. 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	<p>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.</p> <ul style="list-style-type: none">■ System Degradation■ Change■ Environment■ Notification■ Data Availability■ Collector Down■ Object Error

Table 4-94. Symptoms Workspace Options for Message Events (continued)

Option	Description
Symptom definition workspace	Drag the event type to the right pane.
Message Event	<p>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.</p> <p>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.</p> <p>Configure the options:</p> <ul style="list-style-type: none"> ■ 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. ■ 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.

Fault Symptoms

Fault symptoms are based on events published by monitored systems. vRealize Operations Manager correlates a subset of these events and delivers them as faults. Faults are intended to signify events in the monitored systems that affect the availability of objects in your environment. You define symptoms based on faults to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

You can create fault symptoms for the supported published faults. Some object types have multiple fault definitions from which to choose, while others have no fault definitions.

If the adapter published fault definitions for an object type, you can select one or more fault events for a given fault while you define the symptom. The symptom is triggered if the fault is active because of any of the chosen events. If you do not select a fault event, the symptom is triggered if the fault is active because of a fault event.

Fault Symptom Definitions

The Fault Symptom Definitions is a list of the fault-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined fault message events and to determine whether to add, edit, or clone symptoms.

Where You Find Fault Symptoms

To manage symptoms based on fault message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Fault** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-95. Fault Symptoms Definitions Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.</p> <ul style="list-style-type: none"> ■ Add. Add a symptom definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge. ■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge. ■ Clone. Create a copy of the selected symptom definition. ■ Export and Import. Export the file as xml from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to symptoms matching the filter.
Symptom	Descriptive name of the symptom.
Adapter Type	Adapter type for which the symptom is configured.
Object Type	Base object type against which the symptom is defined.
Fault	Selected fault based on object type.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Fault Symptoms Definition Workspace

You define fault symptoms, which are based on events published by the monitored systems, so that you can add one or more symptoms to an alert definition. You use the triggered symptoms to resolve alerts or troubleshoot other problems in vRealize Operations Manager .

How Fault Symptom Definitions Work

A fault symptom is triggered when a fault is active on the base object because of the occurrence of any of the fault events selected in the symptom definition.

Where You Find the Fault Symptom Definition Workspace

To define symptoms based on fault message events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Fault** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-96. 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	<p>Select the fault definition for the selected base object type.</p> <p>Some object types do not have fault definitions, and other types have multiple definitions.</p>
Symptom definition workspace	Drag the fault definition to the right pane.
Fault symptom definition	<p>The fault events are published events from monitored systems. You can configure a single fault event symptom or add multiple symptoms.</p> <p>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:</p> <ul style="list-style-type: none"> ■ Fault event. Select one or more fault events that activate the fault. If you do not select a string, then any of the provided strings are evaluated. ■ Criticality level. Severity of the symptom when it is triggered. ■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms. ■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true. ■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.

Metric Event Symptoms

Metric event symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not vRealize Operations Manager .

Metric event symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that vRealize Operations Manager is actively monitoring.

The metric event thresholds, which determine whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system, represent the type and subtype combination that is specified in the incoming metric event.

- Above Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_ABOVE` defined in the vRealize Operations Manager API SDK.
- Below Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_BELOW` defined in the vRealize Operations Manager API SDK.
- Equal Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_EQUAL` defined in the vRealize Operations Manager API SDK.
- Not Equal Threshold. Corresponds to type and subtype constants `EVENT_CLASS_HT` and `EVENT_SUBCLASS_NOT_EQUAL` defined in the vRealize Operations Manager API SDK.

Metric Event Symptom Definitions

The Metric Event Symptom Definitions is a list of the metric event-based symptoms defined in your vRealize Operations Manager environment. You use the information in the list to evaluate the defined threshold triggering states for the metric events and to determine if you want to add, edit, or clone symptoms.

Where You Find Metric Event Symptoms

To manage symptoms based on metric events, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Click the **Metric Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-97. Metric Event Symptom Definitions Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.</p> <ul style="list-style-type: none"> ■ Add. Add a symptom definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ 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 object type.
Event Type	Specifies whether the metric was above, below, equal to, or not equal to the threshold set by the monitoring system.
Defined By	Indicates whether the symptom was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the symptom was last modified.
Modified By	Displays the name of the user who last modified the symptom.

Metric Event Symptoms Definition Workspace

You define metric event symptoms, which are based on reported violations of metric thresholds from monitored systems, so that you can create one or more of the symptoms that you can add to an alert definition in vRealize Operations Manager .

How Metric Event Symptom Definitions Work

A metric event symptom is triggered when vRealize Operations Manager receives a metric event for the metric and event type defined in the symptom. The event type specifies whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system.

Where You Find the Metric Event Symptom Definition Workspace

To define symptoms based on metric events, in the left pane, in the menu, click **Alerts** and then in the left pane, click **Configuration > Symptom Definitions**. Select the **Metric Event** tab and click **Add** to define a property-based symptom in the workspace.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

Table 4-98. 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 4-98. Symptoms Workspace Options for Metric Events (continued)

Option	Description
Symptom definition workspace	Click and drag the metric to the right pane.
Metric Event	<p>You can configure a single threshold or add multiple thresholds.</p> <p>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.</p> <p>Configure the options:</p> <ul style="list-style-type: none"> ■ 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 4-99. Negative Symptoms Effect on Generated Alert Criticality

Alert Definition Criticality	Negative Symptom Configured Criticality	Standard Symptom Configured Criticality	Alert Criticality When Triggered
Warning	One Critical Symptom	One Immediate Symptom	Warning. The alert criticality is based on the defined alert criticality.
Symptom Based	One Critical Symptom	One Warning Symptom	Warning. The negative symptom has no associated criticality and the criticality of the standard symptom determines the criticality of the generated alert.
Symptom Based	One Critical Symptom	No standard symptom included	Info. Because an alert must have a criticality and the negative alert does not have an associated criticality, the generated alert has a criticality of Info, which is the lowest possible criticality level.

Defining Recommendations for Alert Definitions

Recommendations are instructions to your users who are responsible for responding to alerts. You add recommendations to vRealize Operations Manager alerts so that your users can maintain the objects in your environment at the required levels of performance.

Recommendations provide your network engineers or virtual infrastructure administrators with information to resolve alerts.

Depending on the knowledge level of your users, you can provide more or less information, including the following options, in any combination.

- One line of instruction.
- Steps to resolve the alert on the target object.
- Hyperlink to a Web site, runbook, wiki, or other source.
- Action that makes a change on the target object.

When you define an alert, provide as many relevant action recommendations as possible. If more than one recommendation is available, arrange them in priority order so that the solution with the lowest effect and highest effectiveness is listed first. If no action recommendation is available, add text recommendations. Be as precise as possible when describing what the administrator should do to fix the alert.

Recommendations

Recommendations are probable solutions for an alert generated in vRealize Operations Manager . You can create a library of recommendations that include instructions to your environment administrators or actions that they can run to resolve an alert.

Where You Find Recommendations

To define recommendations, in the menu, click **Alerts** and then in the left pane, click **Configuration > Recommendations**.

You can also define recommendations when you create an alert definition.

Table 4-100. Recommendations Overview Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your recommendations.</p> <ul style="list-style-type: none"> ■ Add. Add a recommendation. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Modify the selected recommendation. ■ Delete. Remove the selected recommendation. ■ Clone. Create a copy of the selected recommendation so that you can create a new recommendation that uses the current one. ■ Export and Import. Export the file as XML from one vRealize Operations Manager so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file.
Filter options	Limits the list to recommendations matching the filter.
Description	Recommendation text as it appears when the alert is generated and the recommendation is presented.
Action	If the recommendation includes running an action, the name of the actions.
Alert Definitions	Displays the number of alert definitions assigned for a particular recommendation. Click this link to view the alert definitions assigned for a particular recommendation and click Remove from all to remove the selected recommendation from all alert definitions.
Defined By	Indicates whether the recommendation was created by a user or provided with a solution adapter.
Last Modified	Displays the date on which the recommendation was last modified.
Modified By	Displays the name of the user who last modified the recommendation.

Recommendation Workspace

You create recommendations that are solutions to alerts generated in vRealize Operations Manager . The recommendations are intended to ensure that your network operations engineers and virtual infrastructure administrators can respond to alerts as quickly and accurately as possible.

How the Recommendations Workspace Works

A recommendation is instructions to your users or actions that your users can perform to resolve an alert. The instructions can be links to useful Web sites or local runbooks, instructions as text, or actions that you can initiate from vRealize Operations Manager .

Where You Find Recommendations Workspace

To define recommendations, click **Alerts** and then in the left pane, click **Configuration > Recommendations**. Click **Add** to create a recommendation.

You can also define recommendations when you define alerts.

Table 4-101. Define Recommendation Options

Option	Description
Create a hyperlink	Enter text in the text box, select the text, and click the button to make the text a hyperlink to a Web site or local wiki page. You cannot modify a hyperlink. To change the link, delete the hyperlinked word and create a new link.
Enter text	Enter the description of what must be done to resolve the triggered alert. The description can include steps a user must take to resolve the alert or it might be instructions to notify a virtual infrastructure administrator. This is a text field.
Adapter Type	Select an adapter type from the drop-down list to narrow down the list of actions displayed in the Actions field.
Action	You can add an action as a method to resolve a triggered symptom or a generated alert. Actions must already be configured in vRealize Operations Manager . You must provide text in the text box to describe the action before you can save the recommendation.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- Set Memory for VM Power Off Allowed
- Set CPU Count for VM Power Off Allowed

■ Set CPU Count and Memory for VM Power Off Allowed

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

Alert Definitions

Alert definitions are a combination of symptoms and recommendations that you combine to identify problem areas in your environment and generate alerts on which you can act for those areas. You use the Alert Definitions to manage your vRealize Operations Manager alert library, and to add or modify the definitions.

Where You Find Alert Definitions

To manage your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**.

Table 4-102. Alert Definition Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your alert definitions.</p> <ul style="list-style-type: none"> ■ Add. Add an alert definition. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Modify the selected definition. ■ Delete. Remove the selected definition. ■ Clone. Create a copy of the selected definition so that you can customize it for your needs. ■ Export and Import. Export the selected definition so that you can import it on another vRealize Operations Manager instance.
Filtering options	<p>Limits the list of alerts to those matching the filter you create.</p> <p>You can also sort on the columns in the data grid.</p>
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	<p>Metadata that is used to classify the alert when it is generated.</p> <p>You define the value on the Alert Impact page of the workspace.</p>
Alert Subtype	<p>Subcategory of the alert type and is the metadata that is used to classify the alert when it is generated.</p> <p>You define the value on the Alert Impact page of the workspace.</p>

Table 4-102. Alert Definition Options (continued)

Option	Description
Criticality	Severity of the alert when it is generated. The criticality includes the following possible values: <ul style="list-style-type: none"> ■ Symptom. Alert is configured to display symptom based criticality. ■ Critical ■ Immediate ■ Warning ■ Info
Impact	Alert is configured to affect the Health, Risk, or Efficiency badge.
Defined by	Indicates who added the alert definition. The alert can be added by an adapter, a user, or the vRealize Operations Manager system.
Last Modified	Displays the date on which the alert was last modified.

Alert Definition Workspace

The alert definition process includes adding symptoms that trigger an alert and recommendations that help you resolve the alert. The alert definitions you create with this process are saved to your vRealize Operations Manager Alert Definition Overview list and actively evaluated in your environment based on your configured policies.

How the Alert Definition Workspace Works

You use the workspace to build alert definitions as you create the definition, the name, description, base object, and the alert impact. You can create or reuse existing symptoms and recommendations as part of the alert definition. If you create symptoms and recommendations, you add them to the definition, and they are added to the symptom and recommendations content libraries for future use. You also enable policies and select notifications for the alerts.

Where You Create an Alert Definition

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**. Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition.

Alert Definition Workspace Options

An alert definition is identified by a name and description. The definition comprises a target object type that is monitored for the alert, the badge that the alert affects, the set symptoms that trigger the alert, the recommendations that might resolve the alert, the policies that are enabled for an alert and the notification setting for which you want to receive the alert.

■ [Alert Definition Workspace Add Alert Details](#)

The name, description, base object type, and other details of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager .

■ [Alert Definition Workspace Add Symptom Definitions](#)

The add symptom definitions options are the mechanisms you use to add existing symptoms or to create new symptoms for the alert definition. If the symptom that you need for an alert definition does not exist, you can create it from this workspace.

■ [Alert Definition Workspace Add Recommendations](#)

Recommendations are instructions you provide to your user so that they can resolve generated alerts. The recommendations might include actions.

■ [Alert Definition Workspace Select Policies](#)

A policy is a set of rules that you define. It allows you to analyze and display information about the objects in your environment.

■ [Alert Definition Workspace Select Notifications](#)

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager .

Alert Definition Workspace Add Alert Details

The name, description, base object type, and other details of the alert definition. This is the information that identifies the alert when it is generated in vRealize Operations Manager .

Where You Define the Alert Details

To create or edit your alert definitions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**. Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition. In the workspace, on the right, enter the details of alert definition.

Table 4-103. Alert Definition Details

Option	Description
Name	Name of the alert as it appears when the alert is generated.
Description	Description of the alert as it appears when the alert is generated. Provide a useful description for your users.
Base Object Type	<p>The object type against which the alert definition is evaluated and the alert is generated.</p> <p>The drop-down menu includes all of the object types in your environment. You can define an alert definition based on one object type.</p>
Impact	<p>Under Advanced Settings, select the badge that is affected if the alert is generated.</p> <p>You can select a badge based on the urgency of the alert.</p> <ul style="list-style-type: none"> ■ 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.

Table 4-103. Alert Definition Details (continued)

Option	Description
Criticality	<p>Severity of the alert that is communicated as part of the alert notification.</p> <p>Select one of the following values.</p> <ul style="list-style-type: none"> ■ 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. If you negate symptoms, the negative symptoms do not contribute to the criticality of a symptom-based alert.
Alert Type and Subtype	<p>Select the type and subtype of alert.</p> <p>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.</p> <p>You can use the type and subtype information to route the alert to the appropriate personnel and department in your organization.</p>
Wait Cycle	<p>The symptoms included in the alert definition remain triggered for this number of collection cycles before the alert is generated.</p> <p>The value must be 1 or greater.</p> <p>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.</p>
Cancel Cycle	<p>The symptoms are cancelled for this number of collection cycles after which the alert is cancelled.</p> <p>The value must be 1 or greater.</p> <p>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.</p>

Click **Next** to add symptom definitions.

Alert Definition Workspace Add Symptom Definitions

The add symptom definitions options are the mechanisms you use to add existing symptoms or to create new symptoms for the alert definition. If the symptom that you need for an alert definition does not exist, you can create it from this workspace.

How the Add Symptom Definitions Options Work

You can select and add symptoms defined for the base object type, and you can add symptoms for related object types. As you add one or more symptoms, you create a symptom expression. If this expression is evaluated as true, then the alert is generated.

Add Symptoms Definitions Options

To add symptom definitions, you can drag the selected symptom in to the left pane. Use the workspace on the left to specify whether all or any of the symptoms or symptom sets must be true to generate an alert.

Table 4-104. Add Symptoms Selection Options

Option	Description
Select Symptom	<p>Select the type of symptom definition that you are adding for the current Defined On object type.</p> <ul style="list-style-type: none"> ■ Metric / Property. Add symptoms that use metric and property symptoms. These metrics are based on the operational or performance values, and configuration properties that vRealize Operations Manager collects from target objects in your environment. ■ Message Event. Add symptoms that use message event symptoms. These symptoms are based on events received as messages from a component of vRealize Operations Manager or from an external monitored system through the system's REST API. ■ Fault Event. Add symptoms that use fault symptoms. These symptoms are based on events that monitored systems publish. vRealize Operations Manager correlates a subset of these events and delivers them as faults. Faults are intended to signify events in the monitored systems that affect the availability of objects in your environment. ■ Metric Event. Add symptoms that use metric event symptoms. These symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not vRealize Operations Manager. These symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that vRealize Operations Manager is actively monitoring. ■ Smart Early Warning. Add a symptom that uses a defined condition that is triggered when the number of anomalies on an object is over the trending threshold. This symptom represents the overall anomalous behavior of the object. Anomalies are based on vRealize Operations Manager analysis of the number of applicable metrics that violate the dynamic threshold that determines the normal operating behavior of the object. This symptom is not configurable. You either use it or you do not use it.
Defined On	<p>Object that the symptom evaluates.</p> <p>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.</p> <ul style="list-style-type: none"> ■ 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	<p>Available only when you select a Defined On value other than Self.</p> <p>Limits the symptoms to those that are configured for the selected object type based on the selected Defined On relationship.</p>

Table 4-104. Add Symptoms Selection Options (continued)

Option	Description
Create New Symptom	<p>If symptoms that you need for your alert do not exist, you can create them.</p> <p>Opens the symptoms definition dialog box.</p> <p>Not available for Smart Early Warning symptoms, which are predefined in the system.</p>
All Filters	<p>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.</p> <ul style="list-style-type: none"> ■ Symptom. Type text to search on the name of the symptom definitions. For example, to display all symptom definitions that have efficiency in their name, type Efficiency. ■ Defined By. Type text to search for the name of the adapter that defines the symptom definitions. For example, to display all symptom definitions provided by the vCenter Adapter, type vCenter. To display only user-defined symptom definitions, type the search term User. <p>To clear a filter, click the double arrow icon that appears next to the filter name.</p>
Quick filter (Name)	Search the list based on the symptom name.
Symptoms list	<p>List of existing symptoms for the selected object type. To configure a symptom, drag it into the left workspace.</p> <p>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.</p>

Use the workspace to configure the interaction of the symptoms and symptom sets.

Table 4-105. Symptom Sets in the Alert Definition Workspace

Option	Description
Trigger alert when {operator} of the symptom sets are true	<p>Select the operator for all of the added symptom sets. Available only when you add more than one symptom set.</p> <ul style="list-style-type: none"> ■ All. All of the symptom sets must be true before the alert is generated. Operates as a Boolean AND. ■ Any. One or more of the symptom sets must be true before the alert is generated. Operates as a Boolean OR.
Symptoms	<p>The symptom sets comprise an expression that is evaluated to determine if an alert should be triggered.</p> <p>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.</p>
Symptom sets	<p>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.</p> <p>A symptom set can include one or more symptoms, and an alert definition can include one or more symptom sets.</p> <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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. <ul style="list-style-type: none"> ■ 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. <p>When you include a symptom in a symptom set, the condition must become true to trigger the symptom set. However, you might want to configure a symptom set where the absence of a symptom condition triggers a symptom. To use the absence of the symptom condition, click the vertical ellipsis on the left of the symptom name and select Invert Symptom.</p> <p>Although you can configure symptom criticality, if you invert a symptom, it does not have an associated criticality that affects the criticality of generated alerts.</p>

Click **Next** to add recommendations.

Alert Definition Workspace Add Recommendations

Recommendations are instructions you provide to your user so that they can resolve generated alerts. The recommendations might include actions.

How Add Recommendations Works

Recommendations are information provided to users to resolve a problem when an alert is generated. You use the recommendation options to add existing information or to create solutions to alerts. If the recommendation that you need for an alert definition does not exist, you can create it from this workspace.

Add Recommendations Options

To add recommendations, you can drag the selected recommendation in to the left pane. Use the workspace on the left to to change the priority order.

Table 4-106. Add Recommendations Options in the Alert Definition Workspace

Option	Description
Create New Recommendation	If recommendations that you need to resolve the symptoms in the problem do not exist, you can create them.
All Filters	<p>Filter the list of recommendations.</p> <ul style="list-style-type: none"> ■ Description. Type text to search on the name of the recommendation. For example, to display all recommendations that have memory in their name, type Memory. ■ Defined By. Type text to search for the name of the adapter that defines the recommendation. For example, to display all recommendations provided by the vCenter Adapter, type vCenter. <p>To clear a filter, click the double arrow icon that appears next to the filter name.</p>
Quick filter (Name)	Limits the list based on the text you enter.
List of available recommendations.	<p>List of existing recommendations that you can drag to the workspace.</p> <p>Recommendations are instructions and, where possible, actions that assist you with resolving alerts when they are triggered.</p>
Recommendation workspace	<p>Add one or more recommendations to the workspace.</p> <p>If you add more than one recommendation, you can drag the recommendations to change the priority order.</p>

Click **Next** to enable policies.

Alert Definition Workspace Select Policies

A policy is a set of rules that you define. It allows you to analyze and display information about the objects in your environment.

How the Select Policies Option Works

Policies define the settings that vRealize Operations Manager applies to your objects when it collects data from your environment. You can select the policies that you want to apply for a particular alert.

Select Policies Option

You can view the policy tree in the left pane and you can either select the default policy or any other policy from the tree. You can also customize thresholds for a policy by clicking the policy and editing the trigger value in the right pane.

Note If you create an alert without enabling any policies, then the alert remains inactive.

Click **Next** to select notifications.

Alert Definition Workspace Select Notifications

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager .

How the Select Notifications Option Works

You can send alert notifications for an alert by assigning the alert to a notification rule that you have set up.

Select Notifications Option

You can view the notification setting on the left pane and select the notification setting for which you want to receive the alert.

Click **Create** to create the alert. The new alert appears in the list of alert definitions.

Create a Simple Alert Definition

While troubleshooting, you can now quickly create an alert for a particular object type or a metric in a quick and efficient way.

You can create a simple alert definition from the following locations.

- In the **Home** page, click **Troubleshoot > Workbench** and select the metric for which you want to create an alert. You can create an alert from the **Potential Evidence** or the **Metrics** tab.
- In the **Alerts** page, click **Triggered Alerts**. Select an alert and click the **Potential Evidence** tab.

Procedure

- 1 Click the drop-down menu available in the right side of the widget and select the **Create an Alert Definition** option.
- 2 In the Create Alert Definition page, enter the **Name** and **Description** of the alert.
- 3 Set thresholds, criticality, and the number of wait cycles. Click **Show Advanced Settings** to set Wait Cycle and Cancel Cycle.

Note The Object Type or Metric/Property are pre-selected and cannot be edited.

4 Click **Create**.

The new alert is created and the policy the object belongs to and its children policies are enabled for the alert.

Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for vRealize Operations Manager to alert you. When the alert is triggered on your host system, vRealize Operations Manager alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have vRealize Operations Manager notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, select **Configuration > Alert Definitions**.

- 2 Enter **capacity** in the search text box.

Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.

- 3 Click the plus sign to create a new capacity alert definition for your host systems.

- a In the alert definition workspace, for the Name and Description, enter **Hosts - Alert on Capacity Exceeded**.
- b For the Base Object Type, select **vCenter Adapter > Host System**
- c For the Alert Impact, select the following options.

Option	Selection
Impact	Select Risk .
Criticality	Select Immediate .
Alert Type and Subtype	Select Application : Capacity .
Wait Cycle	Select 1 .
Cancel Cycle	Select 1 .

- d For Add Symptom Definitions, select the following options.

Option	Selection
Defined On	Select Self .
Symptom Definition Type	Select Metric / Supermetric .
Quick filter (Name)	Enter capacity .

- e From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the right pane.

In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.

- f For Add Recommendations, enter **virtual machine** in the quick filter text box.
- g Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the right pane.

This recommendation is set to Priority 1.

- 4 Click **Save** to save the alert definition.

Your new alert appears in the list of alert definitions.

Results

You have added an alert definition to have vRealize Operations Manager alert you when the capacity of your host systems begins to run out.

Alert Definition Best Practices

As you create alert definitions for your environment, apply consistent best practices so that you optimize alert behavior for your monitored objects.

Alert Definitions Naming and Description

The alert definition name is the short name that appears in the following places:

- In data grids when alerts are generated
- In outbound alert notifications, including the email notifications that are sent when outbound alerts and notifications are configured in your environment

Ensure that you provide an informative name that clearly states the reported problem. Your users can evaluate alerts based on the alert definition name.

The alert definition description is the text that appears in the alert definition details and the outbound alerts. Ensure that you provide a useful description that helps your users understand the problem that generated the alert.

Wait and Cancel Cycle

The wait cycle setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition goes into effect after the wait cycle for the symptom definition results in a triggered symptom. In most alert definitions you configure the sensitivity at the symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that the alert is immediately generated after all of the symptoms are triggered at the desired symptom sensitivity level.

The cancel cycle setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition goes into affect after the cancel cycle for the symptom definition results in a cancelled symptom. In most definitions you configure the sensitivity at the symptom level and configure the cancel cycle of alert definition to 1. This configuration ensures that the alert is immediately cancelled after all of the symptoms conditions disappear after the desired symptom cancel cycle.

Create Alert Definitions to Generate the Fewest Alerts

You can control the size of your alert list and make it easier to manage. When an alert is about a general problem that can be triggered on a large number of objects, configure its definition so that the alert is generated on a higher level object in the hierarchy rather than on individual objects.

As you add symptoms to your alert definition, do not overcrowd a single alert definition with secondary symptoms. Keep the combination of symptoms as simple and straightforward as possible.

You can also use a series of symptom definitions to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of Warning while `Volume reached capacity limit` might have a severity level of Critical. The first symptom is not an immediate threat, but the second one is an immediate threat. You can then include the Warning and Critical symptom definitions in a single alert definition with an Any condition and set the alert criticality to be Symptom Based. These settings cause the alert to be generated with the right criticality if either of the symptoms is triggered.

Avoid Overlapping and Gaps Between Alerts

Overlaps result in two or more alerts being generated for the same underlying condition. Gaps occur when an unresolved alert with lower severity is canceled, but a related alert with a higher severity cannot be triggered.

A gap occurs in a situation where the value is $\leq 50\%$ in one alert definition and $\geq 75\%$ in a second alert definition. The gap occurs because when the percentage of volumes with high use falls between 50 percent and 75 percent, the first problem cancels but the second does not generate an alert. This situation is problematic because no alert definitions are active to cover the gap.

Actionable Recommendations

If you provide text instructions to your users that help them resolve a problem identified by an alert definition, precisely describe how the engineer or administrator should fix the problem to resolve the alert.

To support the instructions, add a link to a wiki, runbook, or other sources of information, and add actions that you run from vRealize Operations Manager on the target systems.

Creating and Managing vRealize Operations Manager Alert Notifications

When alerts are generated in vRealize Operations Manager , they appear in the alert details and object details, but you can also configure vRealize Operations Manager to send your alerts to outside applications using one or more outbound alert options.

You configure notification options to specify which alerts are sent out for the Standard Email, REST, SNMP, and Log File outbound alert plug-ins. For the other plug-in types, all the alerts are sent when the target outbound alert plug-in is enabled.

The most common outbound alert plug-in is the Standard Email plug-in. You configure the Standard Email plug-in to send notifications to one or more users when an alert is generated that meets the criteria you specify in the notification settings.

Outbound Settings

You use the Outbound Settings to manage your communication settings so that you can send information to users or applications outside of vRealize Operations Manager .

How Outbound Settings Work

You manage your outbound options from this page, including adding or editing outbound plug-ins, and turning the configured plug-ins on or off. When enabled, the plug-in sends a message to users as email notifications, or sends a message to other applications.

Where You Find Outbound Settings

To manage your outbound settings, select **Administration** in the left pane, and click **Management > Outbound Settings**.

Table 4-107. Outbound Settings Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your Outbound Plug-Ins.</p> <ul style="list-style-type: none"> ■ Add. Opens the Outbound Plug-In dialog box where you configure the connection options for the instance. Select an existing plugin and click the vertical ellipsis to perform the following actions. ■ Edit. Modify the Outbound Plug-In instance details. ■ Delete. Removes the selected plug-in instance. ■ Enable or Disable. Starts or stops the plug-in instance. Disabling an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.
Instance Name	Name that you assigned when you created the plug-in instance.
Plug-In Type	<p>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.</p> <p>The most common plug-in types include standard email, SNMP trap, log file, and REST.</p>
Status	Specifies whether the plug-in is currently running.

Outbound Plug-Ins

Outbound plug-in settings determine how the supported external notification systems connect to their target systems. You configure one or more instances of one or more plug-in types so that you can send data about generated notifications outside of vRealize Operations Manager .

How Outbound Plug-Ins Work

You configure each plug-in with the required information, including destination locations, hosts, ports, user names, passwords, instance name, or other information that is required to send notifications to those target systems. The target systems can include email recipients, log files, or other management products.

Some plug-ins are included with vRealize Operations Manager , and others might be added when you add a management pack as a solution.

Where You Configure Outbound Settings

To add or edit an outbound plug-in, select **Administration** in the top pane, and click **Management > Outbound Settings**. Click **Add** to add a plug-in instance or select a plug-in, click the vertical ellipsis and select **Edit** to edit the existing plug-in.

Outbound Plug-In Configuration Options

The configuration options vary depending on which plug-in you select from the **Plug-In Type** drop-down menu.

To add outbound notification plug-in, see [Add Outbound Notification Plug-Ins in vRealize Operations Manager](#) .

List of Outbound plug-ins in vRealize Operations Manager

vRealize Operations Manager provides outbound plug-ins. This list includes the name of the plug-in and whether you can filter the outbound data based on your notification settings.

If the plug-in supports configuring notification rules, then you can filter the messages before they are sent to the target system. If the plug-in does not support notifications, all messages are sent to the target system, and you can process them in that application.

If you installed other solutions that include other plug-in options, they appear as a plug-in option with the other plug-ins.

Messages and alerts are sent only when the plug-in is enabled.

Table 4-108. Notification Support for Outbound plug-ins

Outbound plug-in	Configure Notification Rules
Automated Action plug-in	No The Automated Action plug-in is enabled by default. If automated actions stop working, select the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only have to provide the instance name.
Log File plug-in	Yes To filter the log file alerts, you can either configure the file named <code>TextFilter.xml</code> or configure the notification rules.

Table 4-108. Notification Support for Outbound plug-ins (continued)

Outbound plug-in	Configure Notification Rules
Smarts SAM Notification plug-in	No
REST Notification plug-in	Yes
Network Share plug-in	No
Standard Email plug-in	Yes
SNMP Trap plug-in	Yes
Slack plug-in	Yes
Service-Now Notification plug-in	Yes

Add Outbound Notification Plug-Ins in vRealize Operations Manager

You add outbound plug-in instances so that you can notify users about alerts or capture alert data outside of vRealize Operations Manager .

You can configure one or more instances of the same plug-in type if you need to direct alert information to multiple target systems.

The Automated Action plug-in is enabled by default. If automated actions stop working, check the Automated Action plug-in and enable it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.

- [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#)

- [Add a REST plug-in for vRealize Operations Manager Outbound Alerts](#)

- [Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.

- [Add a Network Share Plug-In for vRealize Operations Manager Reports](#)

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports only SMB version 2.1.

- [Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

- [Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts](#)

You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

- [Add a Service-Now Notification Plug-In for Outbound Alerts](#)

You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with vRealize Operations Manager . Service Now creates an incident whenever an alert is triggered in vRealize Operations Manager .

- [Notifications - Add a Slack Plug-In for Outbound Notifications](#)

You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run vRealize Operations Manager using alert links to troubleshoot further.

- [Sample Email Alert](#)

Here is a sample email for a newly created alert.

Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts

You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email vRealize Operations Manager alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals.

Prerequisites

Ensure that you have an email user account that you can use as the connection account for the alert notifications. If you choose to require authentication, you must also know the password for this account.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click **Add**.
- 3 From the **Plug-In Type** drop-down menu, select **Standard Email Plugin**.

The dialog box expands to include your SMTP settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

5 Configure the SMTP options appropriate for your environment.

Option	Description
Use Secure Connection	Enables secure communication encryption using SSL/TLS. If you select this option, you must select a method in the Secure Connection Type drop-down menu.
Requires Authentication	Enables authentication on the email user account that you use to configure this SMTP instance. If you select this option, you must provide a password for the user account.
SMTP Host	URL or IP address of your email host server.
SMTP Port	Default port SMTP uses to connect with the server.
Secure Connection Type	Select either SSL/TLS as the communication encryption method used in your environment from the drop-down menu. You must select a connection type if you select Use Secure Connection.
User Name	Email user account that is used to connect to the email server.
Password	Password for the connection user account. A password is required if you select Requires Authentication.
Sender Email Address	Email address that appears on the notification message.
Sender Name	Displayed name for the sender email address.
Receiver Email Address	Receiver's email address.

6 Click **Save**.

7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the Standard Email Plug-In for outbound SMTP alerts is configured and running.

What to do next

Create notification rules that use the Standard Email Plug-In to send a message to your users about alerts requiring their attention. See [User Scenario: Create a vRealize Operations Manager Email Alert Notification](#).

Add a REST plug-in for vRealize Operations Manager Outbound Alerts

You add a REST plug-in so that you can send vRealize Operations Manager alerts to another REST-enabled application where you built a REST Web service to accept these messages.

The REST Plug-In supports enabling an integration, it does not provide an integration. Depending on your target application, you might need an intermediary REST service or some other mechanism that will correlate the alert and object identifiers included in the REST alert output with the identifiers in your target application.

Determine which content type you are delivering to your target application. If you select application/json, the body of the POST or PUT calls that are sent have the following format. Sample data is included.

```
{
  "startDate":1369757346267,
  "criticality":"ALERT_CRITICALITY_LEVEL_WARNING",
  "Risk":4.0,
  "resourceId":"sample-object-uuid",
  "alertId":"sample-alert-uuid",
  "status":"ACTIVE",
  "subType":"ALERT_SUBTYPE_AVAILABILITY_PROBLEM",
  "cancelDate":1369757346267,
  "resourceKind":"sample-object-type",
  "alertName":"Invalid IP Address for connected Leaf Switch",
  "attributeKeyID":5325,
  "Efficiency":1.0,
  "adapterKind":"sample-adapter-type",
  "Health":1.0,
  "type":"ALERT_TYPE_APPLICATION_PROBLEM",
  "resourceName":"sample-object-name",
  "updateDate":1369757346267,
  "info":"sample-info"
}
```

If you select application/xml, the body of the POST or PUT calls that are sent have the following format:

```
<alert>
  <startDate>1369757346267</startDate>
  <criticality>ALERT_CRITICALITY_LEVEL_WARNING</criticality>
  <Risk>4.0</Risk>
  <resourceId>sample-object-uuid</resourceId>
  <alertId>sample-alert-uuid</alertId>
  <status>ACTIVE</status>
  <subType>ALERT_SUBTYPE_AVAILABILITY_PROBLEM</subType>
  <cancelDate>1369757346267</cancelDate>
  <resourceKind>sample-object-type</resourceKind>
  <alertName>Invalid IP Address for connected Leaf Switch</alertName>
  <attributeKeyId>5325</attributeKeyId>
  <Efficiency>1.0</Efficiency>
  <adapterKind>sample-adapter-type</adapterKind>
  <Health>1.0</Health>
  <type>ALERT_TYPE_APPLICATION_PROBLEM</type>
  <resourceName>sample-object-name</resourceName>
  <updateDate>1369757346267</updateDate>
  <info>sample-info</info>
</alert>
```

Note If the alert is triggered by a non-metric violation, the `attributeKeyID` is omitted from the REST output and is not sent.

If the request is processed as POST, for either JSON or XML, the Web service returns an HTTP status code of 201, which indicates the alert was successfully created at the target. If the request is processed as PUT, the HTTP status code of 202, which indicates the alert was successfully accepted at the target.

Prerequisites

Ensure that you know how and where the alerts sent using the REST plug-in are consumed and processed in your environment, and that you have the appropriate connection information available.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click **Add**.
- 3 From the **Plug-In Type** drop-down menu, select **Rest Notification plug-in**.

The dialog box expands to include your REST settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Rest options appropriate for your environment.

Option	Description
URL	URL to which you are sending the alerts. The URL must support HTTPS. When an alert is sent to the REST Web server, the plug-in appends / {alertID} to the POST or PUT call.
User Name	User account on the target REST system.
Password	User account password.
Content Type	Specify the format for the alert output. <ul style="list-style-type: none"> ■ application/json. Alert data is transmitted using JavaScript Object Notation as human-readable text. ■ application/xml. Alert data is transmitted using XML that is human-readable and machine-readable content.
Certificate thumbprint	Thumbprint for the public certificate for your HTTPS service. Either the SHA1 or SHA256 algorithm can be used.
Connection count	Limits the number of simultaneous alerts that are sent to the target REST server. Use this number to ensure that your REST server is not overwhelmed with requests.

- 6 Click **Save**.
- 7 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the REST plug-in for outbound alerts is configured and running.

What to do next

Create notification rules that use the REST plug-in to send alerts to a REST-enabled application or service in your environment. See [User Scenario: Create a vRealize Operations Manager REST Alert Notification](#).

Add a Log File Plug-In for vRealize Operations Manager Outbound Alerts

You add a Log File plug-in when you want to configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes. If you installed vRealize Operations Manager as a multiple node cluster, each node processes and logs the alerts for the objects that it monitors. Each node logs the alerts for the objects it processes.

All alerts are added to the log file. You can use other applications to filter and manage the logs.

Prerequisites

Ensure that you have write access to the file system path on the target vRealize Operations Manager nodes.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click **Add**.
- 3 From the **Plug-In Type** drop-down menu, select **Log File**.

The dialog box expands to include your log file settings.

- 4 In the **Alert Output Folder** text box, enter the folder name.

If the folder does not exist in the target location, the plug-in creates the folder in the target location. The default target location is: `/usr/lib/vmware-vcops/common/bin/`.

- 5 Click **Save**.
- 6 To start the outbound alert service for this plug-in, select the instance in the list and click **Enable** on the toolbar.

Results

This instance of the log file plug-in is configured and running.

What to do next

When the plug-in is started, the alerts are logged in the file. Verify that the log files are created in the target directory as the alerts are generated, updated, or canceled.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports only SMB version 2.1.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.
- 2 From the toolbar, click **Add**.
- 3 From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

The dialog box expands to include your plug-in instance settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.
Password	The password for the domain user account.
Network share root	<p>The path to the root folder where you want to save the reports. You can specify subfolders for each report when you configure the schedule publication.</p> <p>You must enter an IP address. For example, <code>\\IP_address\ShareRoot</code>. You can use the host name instead of the IP address if the host name is resolved to an IPv4 when accessed from the vRealize Operations Manager host.</p> <p>Note Verify that the root destination folder exists. If the folder is missing, the Network Share plug-in logs an error after 5 unsuccessful attempts.</p>

- 6 Click **Test** to verify the specified paths, credentials, and permissions.
The test might take up to a minute.
- 7 Click **Save**.
The outbound service for this plug-in starts automatically.
- 8 (Optional) To stop an outbound service, select an instance and click **Disable** on the toolbar.

Results

This instance of the Network Share plug-in is configured and running.

What to do next

Create a report schedule and configure it to send reports to your shared folder. See [Schedule Reports Overview](#).

Sample Log File Plug-In Output

Here is a sample log file plug-in output.

```
AlertId :: 9fb52c9c-40f2-46a7-a005-01bf24ab75e6

AlertStatus :: Active

AlertControlState :: Open

AlertGenerateTime :: Wed May 06 06:26:05 UTC 2020 (UTC = 1588746365585)

AlertUpdateTime :: Wed May 06 06:26:05 UTC 2020 (UTC = 1588746365585)

AlertMessage :: 9027

AlertSummaryLink :: https://10.27.82.96/ui/index.action#/object/all/1b852a3c-bbdf-41df-a64d-b40af9673b89/alertsAndSymptoms/alerts/9fb52c9c-40f2-46a7-a005-01bf24ab75e6

AlertType :: Storage - Performance

AlertCriticality :: 4

AffectedResourceId :: 1b852a3c-bbdf-41df-a64d-b40af9673b89

AffectedResourceName :: JNJ_6nodes_Large_HA_4_10.27.83.44

AffectedResourceKind :: VirtualMachine

AffectedResourceParentsNames ::
  VM Entity Status:PoweredOn:all
  DistributedVirtualPortgroup:VM-Network-VLAN-820
  VM Entity Status:PoweredOn:vc_evn-hs1-vc.company.com
  VMFolder:Discovered virtual machine
  HostSystem:evn1-hs1-0808.company.com

AffectedResourceAdapterInstanceResourceName ::
  CompanyAdapter Instance:vc_evn-hs1-vc.company.com

AlertOwner ::

Anomalies ::
  VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 25
  VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 15
  VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44 - [virtualDisk:Aggregate of all instances|
totalWriteLatency_average] - HT above 30.5647619047619 > 30

Health ::
4.0
Risk ::
2.0
Efficiency ::
1.0
KPIFiring ::
```

```

AlertTrigger ::
  Resource                                     Message Info
Alarm Reason      Probability                Prediction Time
VirtualMachine:JNJ_6nodes_Large_HA_4_10.27.83.44    HT above 30.5647619047619 > 30    HT
above            Unable to retrieve value    Unable to retrieve value

AlertRootCause ::
null
AlertRootCauseDetails :: null

AlertName :: Virtual machine disk I/O write latency is high

AlertDescription ::
Virtual machine disk I/O write latency is high

```

Add an SNMP Trap Plug-In for vRealize Operations Manager Outbound Alerts

You add an SNMP Trap plug-in when you want to configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.

You can provide filtering when you define a Notification using an SNMP Trap destination.

Prerequisites

Ensure that you have an SNMP Trap server configured in your environment, and that you know the IP address or host name, port number, and community that it uses.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.
- 2 Click **Outbound Settings** and click **Add**.
- 3 From the **Plug-In Type** drop-down menu, select **SNMP Trap Plugin**.
The dialog box expands to include your SNMP trap settings.
- 4 Enter an **Instance Name**.
- 5 Configure the SNMP trap settings appropriate to your environment.

Option	Description
Destination Host	IP address or fully qualified domain name of the SNMP management system to which you are sending alerts.
Port	Port used to connect to the SNMP management system. Default port is 162.
Community	Text string that allows access to the statistics. SNMP Community strings are used only by devices that support SNMPv1 and SNMPv2c protocol.
Username	User name to configure SNMP trap settings in your environment. If the user name is specified, SNMPv3 is considered as the protocol by the plugin. If left blank, SNMPv2c is considered as the protocol by the plugin.
Authentication Protocol	Authentication algorithms available are SHA-224, SHA-256, SHA-384, SHA-512.

Option	Description
Authentication Password	Authentication password.
Privacy Protocol	Privacy algorithms available are AES192, AES256.
Privacy Password	Privacy password.
Engine ID	<p>Engine ID serves as an identifier for the agent. It is used with a hashing function to generate localized keys for authentication and encryption of SNMP v3 messages.</p> <p>It is mandatory to specify the Engine ID when configuring the SNMP Trap plugin. If you do not add the Engine ID and save the SNMP Trap plugin instance, the field is auto-generated the next time you edit the settings.</p>

6 Click **Test** to validate the connection.

Note The Community and Username options are mutually exclusive. Define either one of them to avoid an error. If you add a user name, you can optionally define the Authentication Protocol and Authentication Password followed by the Privacy Protocol and Privacy Password. The privacy protocol and its password cannot be defined independent of the authentication protocol and its password.

Results

This instance of the SNMP Trap plug-in is configured and running.

What to do next

When the plug-in is added, [Configuring Notifications](#) for receiving the SNMP traps.

Add a Smarts Service Assurance Manager Notification Plug-In for vRealize Operations Manager Outbound Alerts

You add a Smarts SAM Notification plug-in when you want to configure vRealize Operations Manager to send alert notifications to EMC Smarts Server Assurance Manager.

This outbound alert option is useful when you manage the same objects in Server Assurance Manager and in vRealize Operations Manager , and you added the EMC Smarts management pack and configured the solution in vRealize Operations Manager . Although you cannot filter the alerts sent to Service Assurance Manager in vRealize Operations Manager , you can configure the Smarts plug-in to send the alerts to the Smarts Open Integration server. You then configure the Open Integration server to filter the alerts from vRealize Operations Manager , and send only those that pass the filter test to the Smarts Service Assurance Manager service.

Prerequisites

- Verify that you configured the EMC Smarts solution. For documentation regarding EMC Smarts integration, see <https://solutionexchange.vmware.com/store>.
- Ensure that you have the EMC Smarts Broker and Server Assurance Manager instance host name or IP address, user name, and password.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management**.

- 2 Click **Outbound Settings** and click **Add**.

- 3 From the **Plug-In Type** drop-down menu, select **Smarts SAM Notification**.

The dialog box expands to include your Smarts settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Smarts SAM notification settings appropriate for your environment.

Option	Description
Broker	Type the host name or IP address of the EMC Smarts Broker that manages registry for the Server Assurance Manager instance to which you want the notifications sent.
Broker Username	If the Smarts broker is configured as Secure Broker, type the user name for the Broker account.
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.

Add a Service-Now Notification Plug-In for Outbound Alerts

You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with vRealize Operations Manager . Service Now creates an incident whenever an alert is triggered in vRealize Operations Manager .

Using Service-Now Notification Plug-In you can send alert notifications to the Service Now ticketing system to create incidents. The incident includes information like the Caller, Category, Subcategory, Business Service, and other attributes related to alerts.

Prerequisites

Ensure that you have log in credentials for Service-Now.

Ensure that you are assigned with IT Infrastructure Library (ITIL) role in Service Now.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.
- 2 From the toolbar, click **Add** and from the **Plug-In Type** drop-down menu, select **Service-Now Notification Plug-in**.

The dialog box expands to include your plug-in instance settings.

- 3 Enter an **Instance Name**.
- 4 Enter the Service Now URL.
`https://dev22418.service-now.com/`
- 5 Enter the user name and password for Service Now.
- 6 Enter a value for the Connection Count.

The connection count represents the maximum number of open connections allowed per node in vRealize Operations Manager .

- 7 To verify the specified paths, credentials, and permissions, click **Test**.
- 8 Click **Save**.

Results

This instance of the Service-Now Notifications plug-in is configured and running.

What to do next

When the plug-in is added, [Configuring Notifications](#) for creating incidents in Service-Now ticketing system.

Notifications - Add a Slack Plug-In for Outbound Notifications

You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run vRealize Operations Manager using alert links to troubleshoot further.

Prerequisites

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.

- 2 From the toolbar, click **Add** and from the **Plugin Type** drop-down menu, select **Slack Plugin**.

The dialog box expands to include your plug-in instance settings.

- 3 Enter an **Instance Name**.

- 4 Enter a value for the **Connection Count**.

The connection count represents the maximum number of open connections allowed per node in vRealize Operations Manager .

- 5 To verify the specified paths, credentials, and permissions, click **Test**.

- 6 Click **Save**.

Results

This instance of the Slack plugin is configured and running.

What to do next

When the plugin is added, [Configuring Notifications](#) for different slack channels.

Sample Email Alert

Here is a sample email for a newly created alert.

```
Alert Definition Name: Node is experiencing swapping due to memory pressure
Alert Definition Description: Node is experiencing swapping due to memory pressure
Object Name : vRealize Operations Manager Node-vRealize Cluster Node
Object Type : vC-Ops-Node
Alert Impact: risk
Alert State : warning
Alert Type : Application
Alert Sub-Type : Performance
Object Health State: info
Object Risk State: warning
```

```
Object Efficiency State: info
Control State: Open
Symptoms:
SYMPTOM SET - self
```

Symptom Name	Object Name	Object ID	Metric	Message Info
Node swap usage at Warning level	vRealize Operations Manager Node-vRealize Cluster Node	50ec874a-2d7d-4e78-98b1-afb26fd67e58	SwapIWorkload	59.183 > 30.0

```
Recommendations:
Notification Rule Name: rule1
Notification Rule Description:
Alert ID : badc2266-935d-4fb9-8594-e2e71e4866fc
VCOps Server - vRealizeClusterNode

Alert details(link)
```

Configuring Notifications

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside vRealize Operations Manager . You configure notification rules for the supported outbound alerts so that you can filter the alerts that are sent to the selected external system.

You use the notifications list to manage your rules. You then use the notification rules to limit the alerts that are sent to the external system. To use notifications, the supported outbound alert plug-ins must be added and running.

With notification rules, you can limit the data that is sent to the following external systems.

- **Standard Email.** You can create multiple notification rules for various email recipients based on one or more of the filter selections. If you add recipients but do not add filter selections, all the generated alerts are sent to the recipients.
- **REST.** You can create a rule to limit alerts that are sent to the target REST system so that you do not need to implement filtering on that target system.
- **SNMP Trap.** You can configure vRealize Operations Manager to log alerts on an existing SNMP Trap server in your environment.
- **Log File.** You can configure vRealize Operations Manager to log alerts to a file on each of your vRealize Operations Manager nodes.

User Scenario: Create a vRealize Operations Manager Email Alert Notification

As a virtual infrastructure administrator, you need vRealize Operations Manager to send email notifications to your advanced network engineers when critical alerts are generated for mmbhost object, the host for many virtual machines that run transactional applications, where no one has yet taken ownership of the alert.

Prerequisites

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see [Create an Alert Definition for Department Objects](#).
- Ensure that at least one instance of the Standard Email Plug-In is configured and running. See [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**.
- 2 Click **Add** to add a notification rule.
- 3 In the **Name** text box, enter a name similar to **Unclaimed Critical Alerts for mmbhost**.
- 4 In the Method area, select **Standard Email Plug-In** from the drop-down menu, and select the configured instance of the email plug-in.
- 5 Configure the email options.
 - a In the **Recipients** text box, enter the email addresses of the members of your advance engineering team, separating the addresses with a semi-colon (;).
 - b To send a second notification if the alert is still active after a specified amount of time, enter the number of minutes in the **Notify again** text box.
 - c Type number of notifications that are sent to users in the **Max Notifications** text box.
- 6 Set the **Notification Status**, you can either enable or disable a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.
- 7 Configure the scope of filtering criteria.
 - a From the **Scope** drop-down menu, select **Object**.
 - b Click **Select an Object** and enter the name of the object.
In this example, type **mmbhost**.
 - c Locate and select the object in the list, and click **Select**.
- 8 Configure the Notification Trigger.
 - a From the **Notification Trigger** drop-down menu, select **Impact**.
 - b From the adjacent drop-down menu, select **Health**.
- 9 In the Criticality area, click **Critical**.
- 10 Expand the Advanced Filters and from the **Alert States** drop-down menu, select **Open**.
The Open state indicates that no engineer or administrator has taken ownership of the alert.
- 11 Click **Save**.

Results

You created a notification rule that sends an email message to the members of your advance network engineering team when any critical alerts are generated for the mmbhost object and the alert is not claimed by an engineer. This email reminds them to look at the alert, take ownership of it, and work to resolve the triggering symptoms.

What to do next

Respond to alert email notifications. See [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 a REST Web service that accepts these messages. You want only alerts where the virtualization alerts that affect availability alert types go to this outside application. You can then use the provided information to initiate a remediation process in that application to address the problem indicated by the alert.

The notification configuration limits the alerts sent to the outbound alert instance to those matching the notification criteria.

Prerequisites

- Verify that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see [Create an Alert Definition for Department Objects](#).
- Verify that at least one instance of the REST plug-in is configured and running. See [Add a REST plug-in for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**.
- 2 Click **Add** to add a notification rule.
- 3 In the **Name** text box, enter a name similar to **Virtualization Alerts for Availability**.
- 4 In the Method area, select **REST Notification Plugin** from the drop-down menu, and select the configured instance of the email plug-in.
- 5 Set the **Notification Status**, you can either enable or disable a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.
- 6 Configure the Notification Trigger.
 - a From the **Notification Trigger** drop-down menu, select **Alert Type**.
 - b Click **Select an Alert Type/Subtype** and select any alert types or subtypes under **Virtualization/Hypervisor Alerts Availability**.
- 7 In the Criticality area, click **Warning**.
- 8 Expand the Advanced Filters and from the **Alert Status** drop-down menu, select **New**.

The New status indicates that the alert is new to the system and not updated.

9 Click **Save**.

Results

You created a notification rule that sends the alert text to the target REST-enabled system. Only the alerts where the configured alert impact is Virtualization/Hypervisor Availability and where the alert is configured as a warning are sent to the target instance using the REST plug-in.

Notifications

You use the Notifications page to manage your individual alert notification rules. The rules determine which vRealize Operations Manager alerts are sent to the supported target systems.

How Notifications Work

You add, manage, and edit your notification rules from this page. To send notifications to a supported system, you must configure and enable the settings for outbound alerts. The supported outbound notification plug-ins include the Standard Email Plug-In, REST plug-in, SNMP Trap plug-in, and the Log File plug-in.

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances.

Where You Find Notifications

To manage your notifications, in the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**.

Table 4-109. Notifications Options

Option	Description
Toolbar options	<p>Use the toolbar options to manage your notification rules.</p> <ul style="list-style-type: none"> ■ Add. Opens the Add Rule dialog box where you configure the filtering options for the notification rule. <p>Click the horizontal ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Delete. Removes the selected rule. ■ Disable or Enable. Disables or enables the selected rule(s). ■ Export or Import. Export the selected notifications to a ".xml" file so that you can import it on another vRealize Operations Manager instance.
Rule Name	<p>Name you assigned when you created the notification rule. Click the vertical ellipsis to perform the following actions.</p> <ul style="list-style-type: none"> ■ Edit. Allows you to edit the selected rule. ■ Delete. Removes the selected rule. ■ Disable or Enable. Disables or enables the selected rule.
Instance	<p>Name of the configured outbound alert instance for the notification rule.</p> <p>Instances are configured as part of the outbound alerts and can indicate different email servers or sender addresses for alert notifications.</p>
Enabled	Displays if the rule is enabled or not.
Email Address	If the rule is for standard email notifications, the alert recipient email addresses are listed.
Object Name	If the rule specifies a notification for a particular object, the object name is listed.

Table 4-109. Notifications Options (continued)

Option	Description
Children	If the rule specifies a notification for a particular object and selected child objects, the child object types are listed.
Last Modified	Displays the date on which the rule was last modified.
Modified By	Displays the name of the user who last modified the rule.

Notification Rule

Notification rules determine which alerts are sent to the target systems. You configure one or more notification rules to limit the data that vRealize Operations Manager sends to systems or recipients.

How Notification Rules Work

Notification rules are filters that limit the data sent to external systems by using outbound alert plug-ins that are supported, configured, and running. Rather than sending all alerts to all your email recipients, you can use notification rules to send specific alerts. For example, you can send health alerts for virtual machines to one or more of your network operations engineers. You can send critical alerts for selected hosts and clusters to the virtual infrastructure administrator for those objects.

Before you can create and manage notification rules, you must configure the outbound alert plug-in instances.

You can configure one filtering selection, or you can configure as many selections as you need so that vRealize Operations sends only the required data to the target external system.

Where You Find Notification Rules

To manage your notifications, in the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.

Table 4-110. 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. <ul style="list-style-type: none"> ■ Type of plugin. Select one of the configured outbound alert plug-in types: Standard Email, REST, SNMP Trap, Log File, Service-Now, and Slack Plugin. ■ Instance. Select the configured instance for the type of plug-in.

Table 4-110. Notification Rule Configuration Selections (continued)

Selections	Description
Method -Standard Email Plugin	<p>Includes plug-in type and the plug-in instance. If you are configuring notifications for standard email, you can add recipients and associated information.</p> <ul style="list-style-type: none"> ■ Recipients. Enter the email addresses of the individuals to whom you are sending email messages that contain alert notifications. If you are sending to more than one recipient, use a semicolon (;) between addresses. ■ Notify again. Number of minutes between notifications messages for active alerts. Leave the text box empty to send only one message per alert. ■ Max Notifications. Number of times to send the notification for the active alert. Leave the text box empty to send only one message per alert. ■ Delay to notify. Number of minutes to delay before sending a notification when a new alert is generated. For example, if the delay is 10 minutes and a new alert is generated, the notification is not sent for 10 minutes. If the alert is canceled in those 10 minutes, the notification is not sent. The notification delay reduces the number of notifications for alerts that are canceled during that time. ■ Description. Enter the text to include in the email message. For example, Attention Host Management team.
Method - Service-Now Notification Plugin	<p>If you are configuring notifications for a Service-Now notification plug-in, you can add instances and associated information.</p> <ul style="list-style-type: none"> ■ Caller. Enter the name of the person who reported the incident or who is affected by the incident. ■ Category. Specify the category to which the incident belongs. ■ Sub Category. Specify the sub category to which the incident belongs. ■ Business Service. Specify the business service of the incident. ■ Contact Type. Enter the contact type. ■ State. Enter the incident state in digits. ■ Resolution Code. Enter the resolution code for the incident. ■ Resolution notes. Enter the resolution notes for the incident. ■ On hold reason. Enter the reason as to why the incident is on hold. ■ Impact. Set the incident impact in digits. Impact measures the business criticality of the affected service. ■ Urgency. Set urgency for the incident in digits. Urgency defines the number of days taken to resolve an incident. ■ Priority. Enter the priority for the incident. Priority defines the sequence in which the incident must be resolved. ■ Assignment Group. Enter the assignment group for the incident. ■ Assigned To. Enter the details of the person to whom the incident is assigned. ■ Severity. Set the severity for the incident in digits. ■ Upon Approval. Specify the next steps to be taken upon incident approval. ■ Problem. Enter the details of the related problem if it exists. ■ Cause by change. Enter the change request which triggered the incident. ■ Change Request. Enter the details for the related change list if it exists.

Table 4-110. Notification Rule Configuration Selections (continued)

Selections	Description
Method - Slack Plugin	<p>If you are configuring notifications for a Slack plugin, add the Webhook URL of Slack. For example, the Webhook URL is in the format: <code>https://hooks.slack.com/services/T00000000/B00000000/XXXXXXXXXXXXXXXXXXXX</code>.</p> <p>Create and authorize an app within Slack to obtain the Webhook URL. For details on creating and authorizing an app within Slack, refer to the Slack Documentation.</p> <p>Once you have created the notification rule, the alerts are displayed within that particular Slack channel with a link to the alert. Click the link to view the details of the alert in the Object Summary page.</p>
Notification Status	Either enable or disable a notification setting. Disabling a notification will stop the alert notification for that setting and enabling it will activate it again.
Filtering Criteria	Note The Filtering Criteria and Advanced Filter sections are same for all the plugins.
Scope	<p>General object type for which you are filtering the alert notifications.</p> <p>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.</p>
Notification Trigger	<p>Alert type and subtypes, impact, or definition that triggers the alert.</p> <p>After you select the trigger type, you configure the specific selections associated with the trigger type. For example, if you select Alert Definition, you then select the alert definition that limits the data to alerts with this definition. You can select multiple alert definitions as conditions for a notification to trigger.</p>
Criticality	Defined criticality of the alert that results in the data being sent to an external system. For example, if you select Critical, then the data that is sent to the external system must also be labeled as critical.
Advanced Filters	
Alert States	Managed state of the alert, either opened, assigned, or suspended.
Alert Status	Current state of the alert, either canceled, updated, or new.
Collector/Group	Configured collectors in your environment. For example, in an environment where you manage multiple vCenter Server instances, you can select a collector for one instance. If you want to distribute email alert notifications between various groups which use different remote collectors, select Default collector group . This option filters alerts by the target collector group.

Create an Alert Definition for Department Objects

As a virtual infrastructure administrator, you are responsible for the virtual machines and hosts that the accounting department uses. You can create alerts to manage the accounting department objects.

You received several complaints from your users about delays when they are using their accounting applications. Using vRealize Operations Manager, you identified the problem as related to CPU allocations and workloads. To better manage the problem, you create an alert definition with tighter symptom parameters so that you can track the alerts and identify problems before your users encounter further problems.

Using this scenario, you create a monitoring system that monitors your accounting objects and provides timely notifications when problems occur.

Add Description and Base Object to Alert Definition

To create an alert to monitor the CPUs for the accounting department virtual machines and monitor host memory for the hosts on which they operate, you begin by describing the alert.

When you name the alert definition and define alert impact information, you specify how the information about the alert appears in vRealize Operations Manager . The base object is the object around which the alert definition is created. The symptoms can be for the base object and for related objects.

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Configuration > Alert Definitions**.
- 2 Click **Add** to add a definition.
- 3 Type a name and description.

In this scenario, type **Acct VM CPU early warning** as the alert name, which is a quick overview of the problem. The description, which is a detailed overview, should provide information that is as useful as possible. When the alert is generated, this name and description appears in the alert list and in the notification.

- 4 From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Host System**.

This alert is based on host systems because you want an alert that acts as an early warning to possible CPU stress on the virtual machines used in the accounting department. By using host systems as the based object type, you can respond to the alert symptom for the virtual machines with bulk actions rather than responding to an alert for each virtual machine.

- 5 Click **Advanced Settings** and configure the metadata for this alert definition.

- a From the **Impact** drop-down menu, select **Risk**.

This alert indicates a potential problem and requires attention in the near future.

- b From the **Criticality** drop-down menu, select **Immediate**.

As a Risk alert, which is indicative of a future problem, you still want to give it a high criticality so that it is ranked for correct processing. Because it is designed as an early warning, this configuration provides a built-in buffer that makes it an immediate risk rather than a critical risk.

- c From the **Alert Type and Subtype** drop-down menu, select **Performance** under **Virtualization/Hypervisor**.

- d To ensure that the alert is generated during the first collection cycle after the symptoms become true, set the **Wait Cycle** to **1**.
- e To ensure that the an alert is removed as soon as the symptoms are no longer triggered, set the **Cancel Cycle** to **1**.

The alert is canceled in the next collection cycle if the symptoms are no long true.

These alert impact options help you identify and prioritize alerts as they are generated.

Results

You started an alert definition where you provided the name and description, selected host system as the base object type, and defined the data that appears when the alert generated.

What to do next

Continue in the workspace, adding symptoms to your alert definition. See [Add a Virtual Machine CPU Usage Symptom to the Alert Definition](#).

Add a Virtual Machine CPU Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add symptoms to your vRealize Operations Manager alert definition after you provide the basic descriptive information for the alert. The first symptom you add is related to CPU usage on virtual machines. You later use a policy and group to apply alert to the accounting virtual machines.

This scenario has two symptoms, one for the accounting virtual machines and one to monitor the hosts on which the virtual machines operate.

Prerequisites

Begin configuring the alert definition. See [Add Description and Base Object to Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next** and configure the symptoms.
- 2 Begin configuring the symptom set related to virtual machines CPU usage.
 - a From the **Select Symptom** drop-down menu, select **Metric / Property**.
 - b From the **Defined On** drop-down menu, select **Child**.
 - c From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.
 - d Click **Create New** to open the **Add Symptom Definition** workspace window.

- 3 Configure the virtual machine CPU usage symptom in the **Add Symptom Definition** workspace window.

- a From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Virtual Machine**.

The collected metrics for virtual machines appears in the list.

- b In the metrics list **Search** text box, which searches the metric names, type **usage**.
- c In the list, expand **CPU** and drag **Usage (%)** to the workspace on the left.
- d From the threshold drop-down menu, select **Dynamic Threshold**.

Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.

- e In the **Symptom Definition Name** text box, type a name similar to **VM CPU Usage above trend**.
- f From the criticality drop-down menu, select **Warning**.
- g From the threshold drop-down menu, select **Above Threshold**.
- h Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

This Wait Cycle setting requires the symptom condition to be true for 3 collection cycles before the symptom is triggered. This wait avoids triggering the symptom when there is a short spike in CPU usage.

- i Click **Save**.

The dynamic symptom, which identifies when the usage is above the tracked trend, is added to the symptom list.

- 4 In the **Alert Definition Workspace** window, drag **VM CPU Usage above trend** from the symptom definition list to the symptom workspace on the left.

The Child-Virtual Machine symptom set is added to the symptom workspace.

- 5 In the symptoms set, configure the triggering condition so that when the symptom is true on half of the virtual machines in the group to which this alert definition is applied, the symptom set is true.
 - a From the value operator drop-down menu, select **>**.
 - b In the value text box, enter **50**.
 - c From the value type drop-down menu, select **Percent**.

Results

You defined the first symptom set for the alert definition.

What to do next

Add the host memory usage symptom to the alert definition. See [Add a Host Memory Usage Symptom to the Alert Definition](#).

Add a Host Memory Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add a second symptom to your vRealize Operations Manager alert definition after you add the first symptom. The second symptom is related to host memory usage for the hosts on which the accounting virtual machines operate.

Prerequisites

Add the virtual machine CPU usage symptom. See [Add a Virtual Machine CPU Usage Symptom to the Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next**.
- 2 Configure the symptom related to host systems for the virtual machines.
 - a From the **Select Symptom** drop-down menu, select **Metric / Property**.
 - b From the **Defined On** drop-down menu, select **Self**.
 - c Click **Create New** to add new symptom.
- 3 Configure the host system symptom in the **Add Symptom Definition** workspace window.
 - a From the **Base Object Type** drop-down menu, expand **vCenter Adapters** and select **Host System**.
 - b In the metrics list, expand **Memory** and drag **Usage (%)** to the workspace on the left.
 - c From the threshold drop-down menu, select **Dynamic Threshold**.
Dynamic thresholds use vRealize Operations Manager analytics to identify the trend metric values for objects.
 - d In the **Symptom Definition Name** text box, enter a name similar to **Host memory usage above trend**.
 - e From the criticality drop-down menu, select **Warning**.
 - f From the threshold drop-down menu, select **Above Threshold**.
 - g Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.
This Wait Cycle setting requires the symptom condition to be true for three collection cycles before the symptom is triggered. This wait avoids triggering the symptom when a short spike occurs in host memory usage.
 - h Click **Save**.

The dynamic symptom identifies when the hosts on which the accounting virtual machines run are operating above the tracked trend for memory usage.

The dynamic symptom is added to the symptom list.

- 4 In the **Alert Definition Workspace** window, drag **Host memory usage above trend** from the symptoms list to the symptom workspace on the left.

The Self-Host System symptom set is added to the symptom workspace.

- 5 On the Self-Host System symptom set, from the value type drop-down menu for **This Symptom set is true when**, select **Any**.

With this configuration, when any of the hosts running accounting virtual machines exhibit memory usage that is above the analyzed trend, the symptom condition is true.

- 6 At the top of the symptom set list, from the **Match {operator} of the following symptoms** drop-down menu, select **Any**.

With this configuration, if either of the two symptom sets, virtual machine CPU usage or the host memory, are triggered, an alert is generated for the host.

Results

You defined the second symptom set for the alert definition and configured how the two symptom sets are evaluated to determine when the alert is generated.

What to do next

Add recommendations to your alert definition so that you and your engineers know how to resolve the alert when it is generated. See [Add Recommendations to the Alert Definition](#).

Add Recommendations to the Alert Definition

To resolve a generated alert for the accounting department's virtual machines, you provide recommendations so that you or other engineers have the information you need to resolve the alert before your users encounter performance problems.

As part of the alert definition, you add recommendations that include actions that you run from vRealize Operations Manager and instructions for making changes in vCenter Server that resolve the generated alert.

Prerequisites

Add symptoms to your alert definition. See [Add a Host Memory Usage Symptom to the Alert Definition](#).

Procedure

- 1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, **Alert Impact**, and **Add Symptom Definitions**, click **Next** and add the recommended actions and instructions.

- 2 Click **Create New Recommendation** and select an action recommendation to resolve the virtual machine alerts.
 - a In the **Description** text box, enter a description of the action similar to **Add CPUs to virtual machines**.
 - b From the **Actions** drop-down menu, select **Set CPU Count for VM**.
 - c Click **Create**.
- 3 Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory problems similar to this example.

If this host is part of a DRS cluster, check the DRS settings to verify that the load balancing setting are configured correctly. If necessary, manually vMotion the virtual machines.
- 4 Click **Create**.
- 5 Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory alerts.
 - a Enter a description of the recommendation similar to this example.

If this is a standalone host, add more memory to the host.
 - b To make the URL a hyperlink in the instructions, copy the URL, for example, <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html>, to your clipboard.
 - c Highlight the text in the text box and click the hyperlink icon.
 - d Paste the URL in the **Create a hyperlink** text box and click **OK**.
 - e Click **Create**.
- 6 In the **Alert Recommendation Workspace**, drag **Add CPUs to virtual machines**, **If this host is part of a DRS cluster**, and the **If this is a standalone host** recommendations from the list to the recommendation workspace in the order presented.
- 7 Click **Next** to select policies and view notifications.
- 8 Click **Create**.

Results

You provided the recommended actions and instructions to resolve the alert when it is generated. One of the recommendations resolves the virtual machine CPU usage problem and the other resolves the host memory problem.

What to do next

Create a group of objects to use to manage your accounting objects. See [Create a Custom Accounting Department Group](#).

Create a Custom Accounting Department Group

To manage, monitor, and apply policies to the accounting objects as a group, you create a custom object group.

Prerequisites

Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).

Procedure

- 1 In the menu, click **Environment** and click the **Custom Groups** tab.
- 2 Click **Add** to create a new custom group.
- 3 Type a name similar to **Accounting VMs and Hosts**.
- 4 From the **Group Type** drop-down menu, select **Department**.
- 5 From the **Policy** drop-down menu, select **Default Policy**.

When you create a policy, you apply the new policy to the accounting group.

- 6 In the Define membership criteria area, from the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Host System**, and configure the dynamic group criteria.
 - a From the criteria drop-down menu, select **Relationship**.
 - b From the relationships options drop-down menu, select **Parent of**.
 - c From the operator drop-down menu, select **contains**.
 - d In the **Object name** text box, enter **acct**.
 - e From the navigation tree drop-down list, select **vSphere Hosts and Clusters**.

You created a dynamic group where host objects that are the host for virtual machines with acct in the virtual machine name are included in the group. If a virtual machine with acct in the object name is added or moved to a host, the host object is added to the group.

- 7 Click **Preview** in the lower-left corner of the workspace, and verify that the hosts on which your virtual machines that include acct in the object name appear in the **Preview Group** window.
- 8 Click **Close**.
- 9 Click **Add another criteria set**.

A new criteria set is added with the OR operator between the two criteria sets.

- 10 From the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Virtual Machine**, and configure the dynamic group criteria.
 - a From the criteria drop-down menu, select **Properties**.
 - b From the **Pick a property** drop-down menu, expand **Configuration** and double-click **Name**.
 - c From the operator drop-down menu, select **contains**.
 - d In the **Property value** text box, enter **acct**.

You created a dynamic group where virtual machine objects with acct in the object name are included in the group that depends on the presence of those virtual machines. If a virtual machine with acct in the name is added to your environment, it is added to the group.

- 11 Click **Preview** in the lower-left corner of the workspace, and verify that the virtual machines with acct in the object name are added to the list that also includes the host systems.
- 12 Click **Close**.
- 13 Click **OK**.

The Accounting VMs and Hosts group is added to the Groups list.

Results

You created a dynamic object group that changes as virtual machines with acct in their names are added, removed, and moved in your environment.

What to do next

Create a policy that determines how vRealize Operations Manager uses the alert definition to monitor your environment. See [Create a Policy for the Accounting Alert](#).

Create a Policy for the Accounting Alert

To configure how vRealize Operations Manager evaluates the accounting alert definition in your environment, you configure a policy that determines behavior so that you can apply the policy to an object group. The policy limits the application of the alert definition to only the members of the selected object group.

When an alert definition is created, it is added to the default policy and enabled, ensuring that any alert definitions that you create are active in your environment. This alert definition is intended to meet the needs of the accounting department, so you disable it in the default policy and create a new policy to govern how the alert definition is evaluated in your environment, including which accounting virtual machines and related hosts to monitor.

Prerequisites

- Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).

- Verify that you created a group of objects that you use to manage your accounting objects. See [Create a Custom Accounting Department Group](#).

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Policies**.
- 2 Click the **Policy Library** tab and then, click **Add**.
- 3 Type a name similar to **Accounting Objects Alerts Policy** and provide a useful description similar to the following example.

```
This policy is configured to generate alerts when
Accounting VMs and Hosts group objects are above trended
CPU or memory usage.
```

- 4 Select **Default Policy** from the **Start with** drop-down menu.
- 5 On the left, click **Customize Alert / Symptom Definitions** and disable all the alert definitions except the new Acct VM CPU early warning alert.
 - a In the Alert Definitions area, click **Actions** and select **Select All**.
The alerts on the current page are selected.
 - b Click **Actions** and select **Disable**.
The alerts indicate Disabled in the State column.
 - c Repeat the process on each page of the alerts list.
 - d Select **Acct VM CPU early warning** in the list, click **Actions** and select **Enable**.
The Acct VM CPU early warning alert is now enabled.
- 6 On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.
- 7 Click **Save**.

Results

You created a policy where the accounting alert definition exists in a custom policy that is applied only to the virtual machines and hosts for the accounting department.

What to do next

Create an email notification so that you learn about alerts even you when you are not actively monitoring vRealize Operations Manager. See [Configure Notifications for the Department Alert](#).

Configure Notifications for the Department Alert

To receive an email notification when the accounting alert is generated, rather than relying on your ability to generally monitor the accounting department objects in vRealize Operations Manager, you create notification rules.

Creating an email notification when accounting alerts are triggered is an optional process, but it provides you with the alert even when you are not currently working in vRealize Operations Manager .

Prerequisites

- Verify that you completed the alert definition for this scenario. See [Add Recommendations to the Alert Definition](#).
- Verify that standard email outbound alerts are configured in your system. See [Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts](#).

Procedure

- 1 In the menu, click **Alerts** and then in the left pane, click **Configuration > Notifications**.
- 2 Click **Add** to add a notification rule.
- 3 Configure the communication options.
 - a In the **Name** text box, type a name similar to **Acct Dept VMs or Hosts Alerts**.
 - b From the **Select Plug-In Type** drop-down menu, select **StandardEmailPlugin**.
 - c From the **Select Instance** drop-down menu, select the standard email instance that is configured to send messages.
 - d In the **Recipient(s)** text box, type your email address and the addresses of other recipients responsible for the accounting department alerts. Use a semicolon between recipients.
 - e Leave the **Notify again** text box blank.

If you do not provide a value, the email notice is sent only once. This alert is a Risk alert and is intended as an early warning rather than requiring an immediate response.

You configured the name of the notification when it is sent to you and the method that is used to send the message.

- 4 In the Filtering Criteria area, configure the accounting alert notification trigger.
 - a From the **Notification Trigger** drop-down menu, select **Alert Definition**.
 - b Click **Select Alert Definitions**.
 - c Select **Acct VM CPU early warning** and click **Select**.
- 5 Click **Save**.

Results

You created a notification rule that sends you and your designated engineers an email message when this alert is generated for your accounting department alert definition.

What to do next

Create a dashboard with alert-related widgets so that you can monitor alerts for the accounting object group. See [Create a Dashboard to Monitor Department Objects](#).

Create a Dashboard to Monitor Department Objects

To monitor all the alerts related to the accounting department object group, you create a dashboard that includes the alert list and other widgets. The dashboard provides the alert data in a single location for all related objects.

Creating a dashboard to monitor the accounting virtual machines and related hosts is an optional process, but it provides you with a focused view of the accounting object group alerts and objects.

Prerequisites

Create an object group for the accounting department virtual machines and related objects. See [Create a Custom Accounting Department Group](#).

Procedure

- 1 In the menu, click **Dashboards > Actions > Create Dashboard**.
- 2 In the Dashboard Configuration definition area, type a tab name similar to **Accounting VMs and Hosts** and configure the layout options.
- 3 Click **Widget List** and drag the following widgets to the workspace.
 - **Alert List**
 - **Efficiency**
 - **Health**
 - **Risk**
 - **Top Alerts**
 - **Alert Volume**

The blank widgets are added to the workspace. To change the order in which they appear, you can drag them to a different location in the workspace.

- 4 On the Alert List widget title bar, click **Edit Widget** and configure the settings.
 - a In the **Title** text box, change the title to **Acct Dept Alert List**.
 - b For the **Refresh Content** option, select **On**.
 - c Type **Accounting** in the **Search** text box and click **Search**.

The Accounting value corresponds to the name of the object group for the accounting department virtual machines and related hosts.
 - d In the filtered resource list, select the **Accounting VMs and Hosts** group.

The Accounting VMs and Hosts group is identified in the Selected Resource text box.
 - e Click **OK**.

The Acct Dept Alert List is now configured to display alerts for the Accounting VMs and Hosts group objects.

5 Click **Widget Interactions** and configure the following interactions.

- a For Acct Dept Alert List, leave the selected resources blank.
- b For Top Alerts, Health, Risk, Efficiency, and Alert Volume select **Acct Dept Alert List** from the **Selected Resources** drop-down menu.
- c Click **Apply Interactions**.

With the widget interaction configured in this way, the select alert in the Acct Dept Alert List is the source for the data in the other widgets. When you select an alert in the alert list, the Health, Risk, and Efficiency widgets display alerts for that object, Top Alerts displays the topic issues affecting the health of the object, and Alert Volume displays an alert trend chart.

6 Click **Save**.

Results

You created a dashboard that displays the alerts related to the accounting virtual machines and hosts group, including the Risk alert you created.

Alerts Group

For easy and better management of alerts, you can arrange them as a group as per your requirement.


It is complicated to identify a problem in large environments as you receive different kind of alerts. To manage alerts easily, group them by their definitions.

For example, there are 1000 alerts in your system. To identify different types of alerts, group them based on their alert definitions. It is also easy to detect the alert having the highest severity in the group.

When you group alerts, you can see the number of times the alerts having the same alert definition are triggered. By grouping alerts, you can perform the following tasks easily and quickly:

- Find the noisiest alert: The alert that has triggered maximum number of times is known as the noisiest alert. Once you find it, you can disable it to avoid further noise.
- Filter alerts: You can filter alerts based on a substring in alert definitions. The result shows the group of alerts that contain the substring.

Note

- If you cancel or disable an alert group, the alerts are not canceled instantly. It might take some time if the group is large.
 - Only one group can be expanded at a time.
 - The number next to the group denotes the number of alerts in that particular group.
 - The criticality sign  indicates the highest level of severity of an alert in a group.
-

Grouping Alerts

You can group alerts by time, criticality, definition, and object type.

To group alerts:

Procedure

- 1 In the menu, click **Alerts**.
- 2 Select from the various options available from the **Group By** drop-down menu.

Disable Alerts

In an alerts group, you can disable an alert by a single click.

To disable an alert, in the menu, click **Alerts** and then in the left pane, click **Triggered Alerts**. Select the alert name from the data grid, and click **Actions > Disable**.

The alerts can be disabled by two methods:

- **Disable Alert in All Policies:** You disable the alert for all the objects for all the policies.
- **Disable Alert in Selected Policies:** You disable the alert for the objects having the selected policy. Note that this method will work only for objects with alerts.

Configuring Actions

Actions are the ability to update objects or read data about objects in monitored systems, and are commonly provided in vRealize Operations Manager as part of a solution. The actions added by solutions are available from the object Actions menu, list and view menus, including some dashboard widgets, and can be added to alert definition recommendations.

The possible actions include read actions and update actions.

The read actions retrieve data from the target objects.

The update actions modifies the target objects. For example, you can configure an alert definition to notify you when a virtual machine is experiencing memory issues. Add an action in the recommendations that runs the Set Memory for Virtual Machine action. This action increases the memory and resolves the likely cause of the alert.

To see or use the actions for your vCenter Server objects, you must enable actions in the vCenter Adapter for each monitored vCenter Server instance. Actions can only be viewed and accessed if you have the required permissions.

List of vRealize Operations Manager Actions

The list of actions includes the name of the action, the objects that each one modifies, and the object levels at which you can run the action. You use this information to ensure that you correctly apply the actions as alert recommendations and when the actions are available in the **Actions** menu.

Actions and Modified Objects

vRealize Operations Manager actions make changes to objects in your managed vCenter Server instances.

When you grant a user access to actions in vRealize Operations Manager, that user can take the granted action on any object that vRealize Operations Manager manages.

Action Object Levels

The actions are available when you work with different object levels, but they modify only the specified object. If you are working at the cluster level and select **Power On VM**, all the virtual machines in the cluster for which you have access permission are available for you to run the action. If you are working at the virtual machine level, only the selected virtual machine is available.

Table 4-111. vRealize Operations Manager Actions Affected Objects

Action	Modified Object	Object Levels
Rebalance Container	Virtual Machines	<ul style="list-style-type: none"> ■ Data Center ■ Custom Data Center
Delete Idle VM	Virtual Machines	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set DRS Automation	Cluster	<ul style="list-style-type: none"> ■ Clusters
Move VM	Virtual Machine	<ul style="list-style-type: none"> ■ Virtual Machines
Power Off VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Shut Down Guest OS for VM	Virtual Machine VMware Tools must be installed and running on the target virtual machines to run this action.	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Power On VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Delete Powered Off VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set Memory for VM and Set Memory for VM Power Off Allowed	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set Memory Resources for VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines

Table 4-111. 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	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set CPU Resources for VM	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Set CPU Count and Memory for VM and Set CPU Count and Memory for VM Power Off Allowed	Virtual Machine	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Delete Unused Snapshots for VM	Snapshot	<ul style="list-style-type: none"> ■ Clusters ■ Host Systems ■ Virtual Machines
Delete Unused Snapshots for Datastore	Snapshot	<ul style="list-style-type: none"> ■ Clusters ■ Datastores ■ Host Systems
Execute Script	Virtual Machine	<ul style="list-style-type: none"> ■ Virtual Machine
Get Top Processes	Virtual Machine	<ul style="list-style-type: none"> ■ Virtual Machine
Apply Guest User Mapping	vCenter Server	<ul style="list-style-type: none"> ■ vCenter Server <p>Note This action is deprecated and will be removed in the next release.</p>
Clear Guest User Mapping	vCenter Server	<ul style="list-style-type: none"> ■ vCenter Server <p>Note This action is deprecated and will be removed in the next release.</p>
Export Guest User Mapping	vCenter Server	<ul style="list-style-type: none"> ■ vCenter Server <p>Note This action is deprecated and will be removed in the next release.</p>
Configure Included Services	Service Discovery Adapter Instance	<ul style="list-style-type: none"> ■ Service Discovery Adapter Instance <p>Note This action is deprecated and will be removed in the next release.</p>

Actions Overview List in vRealize Operations Manager

Actions are the method you use to configuration changes on managed objects that you initiate from vRealize Operations Manager . These actions are available to add to alert recommendations.

How the Actions Overview List Works

Actions are defined to run on the target object from different object levels, allowing you to add actions as recommendations for alert definitions that are configured for different base objects. The Actions overview is a list of actions available in your environment.

Where You Find the Actions Overview List

To view the available actions, in the menu, click **Alerts** and then in the left pane, click **Configuration > Actions**.

Table 4-112. 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. <ul style="list-style-type: none"> ■ Update actions make changes to the target objects. ■ Read actions retrieve data from the target objects.
Adapter Type	Name of the configured adapter that provides the action.
Resource Adapter Type	Adapter that provides the action.
Associated Object Types	Indicates the object level at which the action instance runs.
Recommendations	Indicates whether the action is used in at least one recommendation.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`
- `Set CPU Count for VM Power Off Allowed`
- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

Actions Supported for Automation

Recommendations can identify ways to remediate problems indicated by an alert. Some of these remediations can be associated with actions defined in your vRealize Operations Manager instance. You can automate several of these remediation actions for an alert when that recommendation is the first priority for that alert.

You enable actionable alerts in your policies. By default, automation is disabled in policies. To configure automation for your policy, in the menu, click **Administration > Policies > Policy Library**. Then, you edit a policy, access the **Alert / Symptom Definitions** workspace, and select **Local** for the **Automate** setting in the Alert / Symptom Definitions pane.

When an action is automated, you can use the **Automated** and **Alert** columns in **Administration > History > Recent Tasks** to identify the automated action and view the results of the action.

- vRealize Operations Manager uses the **automationAdmin** user account to trigger automated actions. For these automated actions that are triggered by alerts, the Submitted By column displays the **automationAdmin** user.
- The Alert column displays the alert that triggered the action. When an alert is triggered that is associated to the recommendation, it triggers the action without any user intervention.

The following actions are supported for automation:

- Delete Powered Off VM
- Delete Idle VM
- Move VM
- Power Off VM
- Power On VM
- Set CPU Count And Memory for VM
- Set CPU Count And Memory for VM Power Off Allowed
- Set CPU Count for VM
- Set CPU Count for VM Power Off Allowed
- Set CPU Resources for VM
- Set Memory for VM
- Set Memory for VM Power Off Allowed
- Set Memory Resources for VM
- Shut Down Guest OS for VM

Roles Needed to Automate Actions

To automate actions, your role must have the following permissions:

- Create, edit, and import policies in **Administration > Policies > Policy Library**.

- Create, clone, edit, and import alert definitions in **Alerts > Configuration > Alert Definitions**.
- Create, edit, and import recommendation definitions in **Alerts > Configuration > Recommendations**.

Important You set the permissions used to run the actions separately from the alert and recommendation definition. Anyone who can modify alerts, recommendations, and policies can also automate the action, even if they do not have permission to run the action.

For example, if you do not have access to the Power Off VM action, but you can create and modify alerts and recommendations, you can see the Power Off VM action and assign it to an alert recommendation. Then, if you automate the action in your policy, vRealize Operations Manager uses the `automationAdmin` user to run the action.

Example Action Supported for Automation

For the Alert Definition named `Virtual machine has chronic high CPU workload leading to CPU stress`, you can automate the action named `Set CPU Count for VM`.

When CPU stress on your virtual machines exceeds a critical, immediate, or warning level, the alert triggers the recommended action without user intervention.

Integration of Actions with vRealize Automation

vRealize Operations Manager restricts actions on objects that vRealize Automation manages, so that the actions do not violate any constraints set forth by vRealize Automation.

When objects in your environment are managed by vRealize Automation, actions in vRealize Operations Manager are not available on those objects. For example, if a host or parent object is being managed by vRealize Automation, actions are not available on that object.

This behavior is true for all actions, including **Power Off VM**, **Move VM**, **Rebalance Container**, and so on.

You cannot turn on or turn off the exclusion of actions on vRealize Automation managed objects.

Actions Determine Whether Objects Are Managed

Actions check the objects in the vRealize Automation managed resource container to determine which objects are being managed by vRealize Automation.

- Actions such as Rebalance Container check the child objects of the data center container or custom data center container to determine whether the objects are managed by vRealize Automation. If the objects are being managed, the action does not appear on those objects.

- The Move VM action checks whether the virtual machine to be moved is being managed by vRealize Automation.

Is the Virtual Machine Managed?	Result of Move VM Action
Yes	The Move VM action does not appear in the vRealize Operations Manager user interface for that virtual machine.
No	The Move VM action moves the virtual machine to a new host, datastore, or new host and datastore. The Move VM action does not check whether the new host or datastore is being managed by vRealize Automation.

- The Delete Snapshots action checks whether the virtual machine or datastore is being managed by vRealize Automation.

Actions on Objects that vRealize Automation Does Not Manage

For a host or parent object that is not managed by vRealize Automation, only the virtual machines that are not being managed by vRealize Automation appear in the action dialog, and you can only take action on the virtual machines that are not being managed by vRealize Automation. If all child objects are being managed by vRealize Automation, the user interface displays the message `No objects are eligible for the selected action.`

If You Attempt to Run an Action on Multiple Objects

If you select multiple objects and attempt to run an action, such as Power Off VM, only the objects that are not being managed by vRealize Automation, which might include a subset of the virtual machines, appear in the Power Off VM action dialog box.

Working with Actions That Use Power Off Allowed

Some of the actions provided with vRealize Operations Manager require the virtual machines to shut down or power off, depending on the configuration of the target machines, to run the actions. You should understand the impact of the Power Off Allowed option before running the actions so that you select the best options for your target virtual machines.

Power Off and Shut Down

The actions that you can run on your vCenter Server instances include actions that shut down virtual machines and actions that power off virtual machines. It also includes actions where the virtual machine must be in a powered off state to complete the action. Whether the VM is shut down or powered off depends on how it is configured and what options you select when you run the action.

The shut-down action shuts down the guest operating system and then powers off the virtual machine. To shut down a virtual machine from vRealize Operations Manager, the VMware Tools must be installed and running on the target objects.

The power off action turns off the VM without regard for the state of the guest operating system. In this case, if the VM is running applications, your user might lose data. After the action is finished, for example, modifying the CPU count, the virtual machine is returned to the power state it was in when the action began.

Power Off Allowed and VMware Tools

For the actions where you are increasing the CPU count or the amount of memory on a VM, some operating systems support the actions if the Hot Plug is configured on the VM. For other operating systems, the virtual machine must be in a powered off state to change the configuration. To accommodate this need where the VMware Tools is not running, the Set CPU Count, Set Memory, and Set CPU Count and Memory actions include the Power Off Allowed option.

If you select Power Off Allowed, and the machine is running, the action verifies whether VMware Tools is installed and running.

- If VMware Tools is installed and running, the virtual machine is shut down before completing the action.
- If VMware Tools is not running or not installed, the virtual machine is powered off without regard for the state of the operating system.

If you do not select Power Off Allowed and you are decreasing the CPU count or memory, or the hot plug is not enabled for increasing the CPU count or memory, the action does not run and the failure is reported in Recent Tasks.

Power Off Allowed When Changing CPU Count or Memory

When you run the actions that change the CPU count and the amount of memory, you must consider several factors to determine if you want to use the Power Off Allowed option. These factors include whether you are increasing or decreasing the CPU or memory and whether the target virtual machines are powered on. If you increase the CPU or memory values, whether hot plug is enabled also affects how you apply the option when you run the action.

How you use Power Off Allowed when you are decreasing the CPU count or the amount of memory depends on the power state of the target virtual machines.

Table 4-113. Decreasing CPU Count and Memory Behavior Based On Options

Virtual Machine Power State	Power Off Allowed Selected	Results
On	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, decreases the CPU or memory, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, decreases the CPU or memory, and powers the machine back on.</p>
On	No	The action does not run on the virtual machine.
Off	Not applicable. The virtual machine is powered off.	The action decreases the value and leaves the virtual machine in a powered off state.

How you use Power Off Allowed when you are increasing the CPU count or the amount of memory depends on several factors, including the state of the target virtual machine and whether hot plug is enabled. Use the following information to determine which scenario applies to your target objects.

If you are increasing the CPU count, you must consider the power state of the virtual machine and whether CPU Hot Plug is enabled when determining whether to apply Power Off Allowed.

Table 4-114. Increasing CPU Count Behavior.

Virtual Machine Power State	CPU Hot Plug Enabled	Power Off Allowed Selected	Results
On	Yes	No	The action increases the CPU count to the specified amount.
On	No	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, increases the CPU count, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, increases the CPU count, and powers the machine back on.</p>
Off	Not applicable. The virtual machine is powered off.	Not required.	The action increases the CPU count to the specified amount.

If you are increasing the memory, you must consider the power state of the virtual machine, whether Memory Hot Plug is enabled, and whether there is a Hot Memory Limit when determining how to apply Power Off Allowed.

Table 4-115. Increasing Memory Amount Behavior

Virtual Machine Power State	Memory Hot Plug Enabled	Hot Memory Limit	Power Off Allowed Selected	Results
On	Yes	New memory value \leq hot memory limit	No	The action increases the memory the specified amount.
On	Yes	New memory value > hot memory limit	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on.</p>
On	No	Not applicable. The hot plug is not enabled.	Yes	<p>If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.</p> <p>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on.</p>
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 the policies in the priority order, as they appear in the priority column. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of any active policy:

- 1 In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

Note The Reorder Policies option is enabled only if there are more than one active policies.

- 2 In the Reorder Policies window, select the policy and drag it up or down to change the priority.
- 3 Click **ok** to save the changes made to the priority.

The priority for the Default Policy is always designated with the letter D, and the other active policies are prioritized with numbers 1, 2, and so on. Policy with priority 1 indicates the highest priority. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Table 4-116. Configurable Policy Rule Elements

Policy Rule Elements	Thresholds, Settings, Definitions
Workload	Configure symptom thresholds for Workload.
Time Remaining	Configure thresholds for the Time Remaining.
Capacity Remaining	Configure thresholds for the Capacity Remaining.
Maintenance Schedule	Sets a time to perform maintenance tasks.
Attributes	An attribute is a collectible data component. You can enable or disable metric, property, and super metric attributes for collection, and set attributes as key performance indicators (KPIs). A KPI is the designation of an attribute that indicates that the attribute is important in your own environment.
Alert Definitions	Enable or disable combinations of symptoms and recommendations to identify a condition that classifies as a problem.
Symptom Definitions	Enable or disable test conditions on properties, metrics, or events.

Privileges to Create, Modify, and Prioritize Policies

You must have privileges to access specific features in the vRealize Operations Manager user interface. The roles associated with your user account determine the features you can access and the actions you can perform. To set the policy priority:

- 1 In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

Note The Reorder Policies option is enabled only if there are more than one active policies.

- 2 In the Reorder Policies window, select the policy and drag it up or down to change the priority.
- 3 Click **ok** to save the changes made to the priority.

How Upgrades Affect Your Policies

After you upgrade vRealize Operations Manager from a previous version, you might find newly added or updated default settings of policies such as, new alerts and symptoms. Hence, you must analyze the settings and modify these settings to optimize them for your current environment. If you apply the policies used with a previous version of vRealize Operations Manager, the manually modified policy settings remain unaltered.

Policy Decisions and Objectives

Implementing policy decisions in vRealize Operations Manager is typically the responsibility of the Infrastructure Administrator or the Virtual Infrastructure Administrator, but users who have privileges can also create and modify policies.

You must be aware of the policies established to analyze and monitor the resources in your IT infrastructure.

- If you are a Network Operations engineer, you must understand how policies affect the data that vRealize Operations Manager reports on objects, and which policies assigned to objects report alerts and issues.
- If you are the person whose role is to recommend an initial setup for policies, you typically edit and configure the policies in vRealize Operations Manager.
- If your primary role is to assess problems that occur in your environment, but you do not have the responsibility to change the policies, you must still understand how the policies applied to objects affect the data that appears in vRealize Operations Manager. For example, you might need to know which policies apply to objects that are associated with particular alerts.
- If you are a typical application user who receives reports from vRealize Operations Manager, you must have a high-level understanding of the operational policies so that you can understand the reported data values.

Policies Library

The policies library displays the base settings, default policy, and other best practice policies that vRealize Operations Manager includes. You can use the policies library to create your own policies. The policies library includes all the configurable settings for the policy elements, such as workload, capacity and time remaining, and so on.

How the Policies Library Works

Use the options in policies library to create your own policy from an existing policy, or to override the settings from an existing policy so that you can apply the new settings to groups of objects. You can also import or export a policy and reorder the policies.

Select a policy to display its details in the right pane. The right pane displays a high-level overview of all the details and options for that policy where these details are categorized in tabs. Expand each category to view all the related details.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for metrics and properties, alerts and symptoms, capacity, compliance, workload automation, and groups and objects. In this workspace, you can also apply the policy to objects and object groups. To update the policy associated with an object or object group, the role assigned to your user account must have the Manage Association permission enabled for policy management.

Where You Manage the Policies Library

To manage the policies library, in the menu, click **Administration**, and then in the left pane click **Policies**. The policies library appears and lists the policies available to use for your environment.

Table 4-117. Policy Library Tab Options

Option	Description
Toolbar	<p>Use the toolbar selections to take action in the policies library.</p> <ul style="list-style-type: none"> ■ Add. Create a policy from an existing policy. ■ Edit. Customize the policy so that you can override settings for vRealize Operations Manager to analyze and report data about the associated objects. ■ Delete. Remove a policy from the list. ■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to D, which gives that policy the highest priority. ■ Import Policy and Export Policy. You can import or export a policy in XML format. To import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management. ■ Reorder Policies. Change the priority of the active policies.
Policies library data grid	<p>vRealize Operations Manager displays the high-level details for the policies.</p> <ul style="list-style-type: none"> ■ Name. Name of the policy as it appears in the Add or Edit Policy workspace, and in areas where the policy applies to objects, such as in Custom Groups. ■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects. ■ Last Modified. Date and time that the policy was last modified. ■ Status. Indicates whether the policy is active or inactive.
Policies library > Right Pane	<p>The right pane displays the name and description of the policy from which the settings are inherited, the policy priority, and the option to edit the policy. From the right pane, you can view the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.</p> <ul style="list-style-type: none"> ■ Metrics and Properties: Displays all the attribute types included in the policy. Attribute type includes, metrics properties, and super metrics. ■ Alerts and Symptoms: Displays all the alert and symptom definitions included in the policy. The Alert Definitions tabs display an overview of the alert definition, criticality, symptom, and state. The Symptoms Definitions tab displays an overview of the symptom name, criticality, and the metric name. ■ Capacity: Displays an overview of all the thresholds of the objects included in the policy. ■ Compliance: Displays the compliance thresholds inherited from the base policy or set while creating the policy. ■ Workload Automation: Displays the details of the workload optimized in your environment per your definition. ■ Groups and Objects: Displays the object or object groups associated with the selected policy and the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it is shown here.

Operational Policies

Determine how to have vRealize Operations Manager monitor your objects, and how to notify you about problems that occur with those objects.

vRealize Operations Manager Administrators assign policies to objects or object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with objects or object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Enable and disable alerts.
- Control data collections by persisting or not persisting metrics on the objects in your environment.
- Configure the product analytics and thresholds.
- Monitor objects and applications at different service levels.
- Prioritize policies so that the most important rules override the defaults.
- Understand the rules that affect the analytics.
- Understand which policies apply to objects or object groups.

vRealize Operations Manager includes a library of built-in active policies that are already defined for your use. vRealize Operations Manager applies these policies in priority order.

When you apply a policy to an object or an object group, vRealize Operations Manager collects data from the objects based on the thresholds, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that are enabled in the policy.

The following examples of policies might exist for a typical IT environment.

- Maintenance: Optimized for ongoing monitoring, with no thresholds or alerts.
- Critical Production: Production environment ready, optimized for performance with sensitive alerting.
- Important Production: Production environment ready, optimized for performance with medium alerting.
- Batch Workloads: Optimized to process jobs.
- Test, Staging, and QA: Less critical settings, fewer alerts.
- Development: Less critical settings, no alerts.
- Low Priority: Ensures efficient use of resources.
- Default Policy: Default system settings.

Types of Policies

There are three types of policies such as default policies, custom policies, and policies that are offered with vRealize Operations Manager.

Custom Policies

You can customize the default policy and base policies included with vRealize Operations Manager for your own environment. You can then apply your custom policy to an individual object or groups of objects, such as the objects in a cluster, or virtual machines and hosts, or to a group that you create to include unique objects and specific criteria.

You must be familiar with the policies so that you can understand the data that appears in the user interface, because policies drive the results that appear in the vRealize Operations Manager dashboards, views, and reports.

To determine how to customize operational policies and apply them to your environment, you must plan ahead. For example:

- Must you track CPU allocation? If you overallocate CPU, what percentage must you apply to your production and test objects?
- Will you overallocate memory or storage? If you use High Availability, what buffers must you use?
- How do you classify your logically defined workloads, such as production clusters, test or development clusters, and clusters used for batch workloads? Or, do you include all clusters in a single workload?
- How do you capture peak use times or spikes in system activity? In some cases, you might need to reduce alerts so that they are meaningful when you apply policies.

When you have privileges applied to your user account through the roles assigned, you can create and modify policies, and apply them to objects. For example:

- Create a policy from an existing base policy, inherit the base policy settings, then override specific settings to analyze and monitor your objects.
- Use policies to analyze and monitor vCenter Server objects and non-vCenter Server objects.
- Set custom thresholds for capacity settings on all object types to have vRealize Operations Manager report on workload, and so on.
- Enable specific attributes for collection, including metrics, properties, and super metrics.
- Enable or disable alert definitions and symptom definitions in your custom policy settings.
- Apply the custom policy to an individual object or groups of objects.

When you use an existing policy to create a custom policy, you override the policy settings to meet your own needs. You set the allocation and demand, the overcommit ratios for CPU and memory, and the thresholds for capacity risk and buffers. To allocate and configure what your environment is actually using, you use the allocation model and the demand model together. Depending on the type of environment you monitor, such as a production environment versus a test or development environment, whether you over allocate at all and by how much depends on the workloads and environment to which the policy applies. You might be more conservative with the level of allocation in your test environment and less conservative in your production environment.

When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

Your policies are unique to your environment. Because policies direct vRealize Operations Manager to monitor the objects in your environment, they are read-only and do not alter the state of your objects. For this reason, you can override the policy settings to fine-tune them until vRealize Operations Manager displays the results that are meaningful and that affect for your environment. For example, you can adjust the capacity buffer settings in your policy, and then view the data that appears in the dashboards to see the effect of the policy settings.

Default Policy in vRealize Operations Manager

The default policy is a set of rules that applies to most of your objects.

The Default policy is marked with the letter D in the Priority column and can apply to any number of objects.

All the Default policies appear in the Default Policy group in the policies library, even if that policy is not associated with an object group. When an object group does not have a policy applied, vRealize Operations Manager associates the Default policy with that group.

A policy can inherit the Default policy settings, and those settings can apply to various objects under several conditions.

The policy that is set to Default always takes the lowest priority. If you attempt to set two policies as the Default policy, the first policy that you set to Default is initially set to the lowest priority. When you set the second policy to Default, that policy then takes the lowest priority, and the earlier policy that you set to Default is set to the second lowest priority.

You can use the Default policy as the base policy to create your own custom policy. You modify the default policy settings to create a policy that meets your analysis and monitoring needs. When you start with the Default policy, your new policy inherits all the settings from the Default base policy. You can then customize your new policy and override these settings.

The data adapters and solutions installed in vRealize Operations Manager provide a collective group of base settings that apply to all objects. In the policy navigation tree in the policies library, these settings are called Base Settings. The Default policy inherits all the base settings by default.

Policies Provided with vRealize Operations Manager

vRealize Operations Manager includes sets of policies that you can use to monitor your environment, or as the starting point to create your own policies.

Verify that you are familiar with the policies provided with vRealize Operations Manager so that you can use them in your own environment, and to include settings in new policies that you create.

Where You Find the Policies Provided with vRealize Operations Manager Policies

In the menu, click **Administration**, and then in the left pane click **Policies** to see the policies provided with vRealize Operations Manager.

Policies That vRealize Operations Manager Includes

All policies exist under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects. In the policies library, these settings are called Base Settings.

The Base Settings policy is the umbrella policy for all other policies, and appears at the top of the policy list in the policies library. All the other policies reside under the Base Settings, because the data adapters and solutions installed in your vRealize Operations Manager instance provide a collective group of base settings that apply to all objects.

The configuration based policy set includes policies provided with vRealize Operations Manager that you use for specific settings on objects to report on your objects. This set includes several types of policies:

- Efficiency alerts policies for infrastructure objects and virtual machines
- Health alerts policies for infrastructure objects
- Overcommit policies for CPU and Memory
- Risk alerts policies for infrastructure objects and virtual machines

The Default Policy includes a set of rules that applies to most of your objects.

Using the Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the policy workspace to create local policies quickly, and update the settings in existing policies. Select a base policy to use as the source for your local policy settings, and modify the thresholds and settings used for analysis and collection of data from objects or object groups in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated objects or object groups.

Prerequisites

Verify that objects or object groups exist for vRealize Operations Manager to analyze and collect data, and if they do not exist, create them. See [Managing Custom Object Groups in vRealize Operations Manager](#) .

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Policies**.
- 2 Click **Add** to add a policy, or select the policy and click **Edit Policy** to edit an existing policy.

You can add and edit policies and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.

- 3 In the Create Policies workspace, assign a name to the policy, and enter the description.

Give the policy a meaningful name and description so that all users know the purpose of the policy.

- 4 From the **Inherit From** drop-down, select one or more policies to use as a baseline to define the settings for your new local policy.

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your new policy settings.

- 5 Click **Create Policy**.

The Create Policies workspace provides the options to customize your policy.

- 6 Click **Metrics and Properties**. In this workspace, select the metric, property, or super metric attributes to include in your policy.

vRealize Operations Manager collects data from the objects in your environment based on the metric, property, or super metric attributes that you include in the policy.

- a Click **Save** and return to the create policies workspace.

- 7 Click **Alerts and Symptoms**. In this workspace, select the alert definitions and symptom definitions, and enable or disable them as required for your policy.

vRealize Operations Manager identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

- a Click **Save** and return to the create policies workspace.

- 8 Click **Capacity**. In this workspace, select and override the situational settings such as committed projects to calculate capacity, time remaining, and other detailed settings.

- a Click **Save** and return to the create policies workspace.

- 9 Click **Compliance**. In this workspace, set the compliance threshold required for your policy.

- a Click **Save** and return to the create policies workspace.

- 10 Click **Workload Automation**. In this workspace, select the optimization settings required for your policy.

Click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

- a Click **Save** and return to the create policies workspace.

- 11 Click **Groups and Objects**. In this workspace, select one or more groups and objects to which the policy applies.

vRealize Operations Manager monitors the objects according to the settings in the policy that is applied to the object or the object group, triggers alerts when thresholds are violated, and reports the results in the dashboards, views, and reports. If you do not assign a policy to one or more objects or object groups, vRealize Operations Manager does not assign the settings in that policy to any objects, and the policy is not active. For an object or an object group that does not have a policy assigned, vRealize Operations Manager associates the object group with the Default Policy.

Filter the object types, and modify the settings for those object types so that vRealize Operations Manager collects and displays the data that you expect in the dashboards and views.

- a Click **Save** and return to the create policies workspace.

What to do next

After vRealize Operations Manager analyzes and collects data from the objects in your environment, review the data in the dashboards and views. If the data is not what you expected, edit your local policy to customize and override the settings until the dashboards display the data that you need.

Policy Workspace in vRealize Operations Manager

The policy workspace allows you to quickly create and modify policies. To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more objects or object groups.

How the Policy Workspace Works

Every policy includes a set of packages, and uses the defined problems, symptoms, metrics, and properties in those packages to apply to specific objects or object groups in your environment. You can view the details for the settings inherited from the base policy, and display specific settings for certain object types. You can override the settings of other policies, and include additional policy settings to apply to the object types.

Use the **Add** and **Edit** options to create policies and edit existing policies.

Where You Create and Modify a Policy

To create and modify policies, in the menu, click **Administration**, and then in the left pane click **Policies** and click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit the policy. The policy workspace is where you select the base policies, and customize and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can apply the policy to objects or object groups.

To remove a policy from the list, select the policy, click the horizontal ellipse, and select **Delete**.

Policy Workspace Options

The policy workspace includes a step-by-step workflow to create and edit a policy, and apply the policy to custom object groups.

- [Getting Started Details](#)

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

- [Select the Inherited Policy Details](#)

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a policy.

- [Capacity Details](#)

You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

- [Compliance Details](#)

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the compliance for the object types in your policy.

- [Workload Automation Details](#)

You can set the workload automation options for your policy, so that vRealize Operations Manager can optimize the workload in your environment as per your definition.

- [Metrics and Properties Details](#)

You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

- [Alert and Symptom Details](#)

You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

- [Groups and Objects details](#)

You can assign your local policy to one or more objects or groups of objects to have vRealize Operations Manager analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

Getting Started Details

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

Where You Assign the Policy Name and Description

To add a name and description to a policy, in the menu, click **Administration**, and then in the left pane click **Policies** and click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. The name and description appear in the Create or Edit policy workspace.

Table 4-118. Name and Description Options in the Create or Edit Policy Workspace

Option	Description
Name	Name of the policy as it appears in the Create or Edit Policy screens, and in areas where the policy applies to objects, such as Custom Groups.
Description	Meaningful description of the policy. For example, use the description to indicate which policy is inherited, and any specific information that users must understand the relationship of the policy to one or more groups of objects.
Inherit From	The base policy that is used as a starting point. All settings from the base policy will be inherited as default settings in your new policy. You can override these settings to customize the new policy. Select a base policy to inherit the policy settings as a starting point for your new policy.

Select the Inherited Policy Details

You can use any of the policies provided with vRealize Operations Manager as a baseline source for your policy settings when you create a policy.

In the policy content area, you can perform the following actions:

- View the packages and elements for the inherited policy and additional policies that you selected to override the settings.
- Compare the differences in settings highlighted between these policies.
- Display object types.

To create a policy, select a base policy to inherit your new custom policy inherits settings. To override some of the settings in the base policy according to the requirements for the service level agreement for your environment, you can select and apply a separate policy for a management pack solution. The override policy includes specific settings defined for the types of objects to override, either manually or that an adapter provides when it is integrated with vRealize Operations Manager. The settings in the override policy overwrite the settings in the base policy that you selected.

When you select and apply a policy to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the policy settings cards.

Click each card to display the inherited policy configuration, and your policy, and displays a preview of the selected policy settings. When you select one of the policy cards, you can view the number of enabled and disabled alert definitions, symptom definitions, metrics and properties, and the number of enabled and disabled changes.

When you select the Groups and Objects card, you select the objects to view so that you can see which policy elements apply to the object type. For example, when you select the StorageArray object type, the workspace displays the local packages for the policy and the object group types with the number of policy elements in each group.

You can preview the policy settings for all object types, only the object types that have settings changed locally, or settings for new object types that you add to the list, such as Storage Array storage devices.

Where You Select and Override Base Policies Settings

To select a base policy to use as a starting point for your own policy, and to select a policy to override one or more settings that your policy inherits from the base policy, in the menu, select **Administration**, and then in the left pane click **Policies** and click **Add** to add a policy. In the Create policies workspace, add a name and description for the policy and from the **Inherit From** drop-down, select the base policy. The policy configuration, objects, and preview appear in cards below this drop-down.

Capacity Details

You can filter the object types, and modify the settings for those object types so that vRealize Operations Manager applies these settings. The data that you expect then appears in the dashboards and views.

How the Capacity Workspace Works

When you turn on and configure the Capacity settings for a policy, you can override the settings for the policy elements that vRealize Operations Manager uses to trigger alerts and display data. These types of settings include symptom thresholds based on alerts, situational settings such as committed projects to calculate capacity and time remaining, and other detailed settings.

Policies focus on objects and object groups. When you configure policy settings for your local policy, you must consider the object type and the results that you expect to see in the dashboards and views. If you do not change these settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

Where You Set the Policy Capacity Settings

To set the capacity settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click the **Capacity** card. The capacity settings for host systems, virtual machines, and other object types that you select appear in the workspace.

You can also edit the capacity settings while working on the objects under the Environment Tab. In the **Capacity** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Capacity Setting**.

Table 4-119. Capacity Settings in the Create or Edit Policy Workspace

Option	Description
Select Object Type	Use the drop-down menu to select object types. Click the All Filters button to add the selected object type to the list so that you can preview and configure the settings. Add settings for a new set of objects. Provide a list of the object types so that you can select an object type, such as Storage Devices > SAN , and add the selected object to the object types list.
All Filters	When you select a filter, a list of the object types that you selected is displayed in the left pane with the threshold settings in the right pane.
Capacity settings for object types	Select an object to view the policy elements and settings for the object type so that you can have vRealize Operations Manager analyze the object type. You can view and modify the threshold settings for the following policy elements: <ul style="list-style-type: none"> ■ Workload ■ Time Remaining ■ Capacity Remaining ■ Maintenance Schedule ■ Allocation Model ■ Custom Profile ■ Capacity Buffer Click the lock icon on the left of each element to override the settings and change the thresholds for your policy.
Time Remaining Calculations	You can set the risk level for the time that is remaining when the forecasted total need of a metric reaches usable capacity. <ul style="list-style-type: none"> ■ Conservative. Select this option for production and mission-critical workloads. ■ Aggressive. Select this option for non-critical workloads.

Policy Workload Element

Workload is a measurement of the demand for resources on an object. You can turn on and configure the settings for the Workload element for the object types in your policy.

How the Workload Element Works

The Workload element determines how vRealize Operations Manager reports on the resources that the selected object group uses. The resources available to the object group depend on the amount of configured and usable resources.

- A specific amount of physical memory is a configured resource for a host system, and a specific number of CPUs is a configured resource for a virtual machine.
- The usable resource for an object or an object group is a subset of, or equal to, the configured amount.
- The configured and usable amount of a resource can vary depending on the type of resource and the amount of virtualization overhead required, such as the memory that an ESX host machine requires to run the host system. When accounting for overhead, the resources required for overhead are not considered to be usable, because of the reservations required for virtual machines or for the high availability buffer.

Where You Override the Policy Workload Element

To view and override the policy workload capacity setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The workload settings for the object types that you have selected appear in the right pane.

View the Workload policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-120. Policy Workload Element Settings in the Create or Edit Policies Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Workload	Allows you to set the number of collection cycles it takes to trigger or clear an alert.

Policy Time Remaining Element

The Time remaining element is a measure of the amount of time left before your objects run out of capacity.

How the Time Remaining Element Works

The Time Remaining element determines how vRealize Operations Manager reports on the available time until capacity runs out for a specific object type group.

- The time remaining indicates the amount of time that remains before the object group consumes the capacity available. vRealize Operations Manager calculates the time remaining as the number of days remaining until all the capacity is consumed.
- To keep the Time Remaining more than the critical threshold setting or to keep it green, your objects must have more days of capacity available.

Where You Override the Policy Time Remaining Element

To view and override the policy Time Remaining capacity setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects. The time remaining settings for the object types that you have selected appear in the right pane.

View the Time Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-121. Policy Time Remaining Element Settings in the Create or Edit Policies Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Time Remaining	Allows you to set the number of days until capacity is projected to run out based on your current consumption trend.

Policy Capacity Remaining Element

Capacity is a measurement of the amount of memory, CPU, and disk space for an object. You can turn on and configure the settings for the Capacity Remaining element for the object types in your policy.

How the Capacity Remaining Element Works

The Capacity Remaining element determines how reports on the available capacity until resources run out for a specific object type group.

- The capacity remaining indicates the capability of your environment to accommodate workload.
- Usable capacity is a measurement of the percentage of capacity available, minus the capacity affected when you use high availability.

Where You Override the Policy Capacity Remaining Element

To view and override the policy Capacity Remaining analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The capacity remaining settings for the object types that you have selected appear in the right pane.

View the Capacity Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-122. Policy Capacity Remaining Element Settings in the Create or Edit Policies Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Capacity Remaining	Allows you to set the percentage at which the capacity remaining alerts must be triggered.

Policy Maintenance Schedule Element

You can set a time to perform maintenance tasks for each policy.

Where You Override the Policy Maintenance Schedule Element

To view and override the policy Maintenance Schedule analysis setting, in the menu, click **Administration**, and then in the left pane, click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The maintenance schedule settings for the object types that you have selected in the workspace appear in the right pane.

View the maintenance schedule policy element.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-123. Policy Maintenance Schedule Element Settings in the Create or Edit Policies Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Maintenance Schedule	Sets a time to perform maintenance tasks. During maintenance, vRealize Operations Manager does not calculate analytics.

Policy Allocation Model Element

Allocation model defines how much CPU, memory, or disk space is allocated to objects in a cluster or datastore cluster. In the policy, you can turn on the Allocation Model element and configure the resource allocation for the objects.

How the Allocation Model Element Works

The Allocation Model element determines how calculates capacity when you allocate a specific amount of CPU, memory, and disk space resource to clusters or datastore clusters. You can specify the allocation ratio for either one, or all of the resource containers of the cluster. Unlike the demand model, the allocation model is used for capacity calculations only when you turn it on in the policy.

The allocation model element also affects the reclaimable resources for memory and storage in Reclaim page. When you turn on the Allocation Model element in the policy, the tabular representation of the VMs and snapshots in the selected data center from which resources can be reclaimed displays reclaimable memory and disk space based on the overcommit values.

Where You Override the Allocation Model Element

To view and override the policy workload analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The allocation model settings for the object types that you selected appear in the right pane.

Click the unlock icon next to Allocation Model to set the overcommit ratios.

Table 4-124. Policy Allocation Model Element Settings

Option	Description
Set overcommit ratio, to enable Allocation Model	Allows you to set the overcommit ratio for CPU, memory, or disk space. Select the check box next to the resource container you want to edit and change the overcommit ratio value.

Policy Custom Profile Element

The custom profile element lets you apply a custom profile which shows how many more of a specified object can fit in your environment depending on the available capacity and object configuration.

Where You Define the Custom Profiles

To define a custom profile, in the menu click **Administration**, and then in the left pane click **Configuration**. Click **Custom Profiles** and click the **Add** option to define a new custom profile.

Where You Select the Custom Profile Element

To view and override the policy Custom Profile analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The custom profile element for the object types that you selected in the workspace appear in the right pane. Click the lock icon to unlock the section and make changes.

Policy Capacity Buffer Element

The capacity buffer element lets you add buffer for capacity and cost calculation. For vCenter Server objects, you can add buffer to CPU, Memory, and Disk Space for the Demand and Allocation models. You can add capacity buffer to clusters and datastore clusters. The values that you define here affect the cluster cost calculation. The time remaining, capacity remaining, and recommended values are calculated based on the buffer. For WLP, capacity buffer is first considered and then the headroom that you have defined is considered.

Where You Define the Capacity Buffer

To view and override the policy Capacity Buffer analysis setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Capacity**, then select one or more objects in the left pane. The custom profile element for the object types that you selected in the workspace appear in the right pane. Click the lock icon to unlock the section and make changes.

How the Capacity Buffer Element Works

The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required. The value of the usable capacity reduces by the buffer amount that you specify here. The default buffer value is zero. If you are upgrading from a previous version of vRealize Operations Manager, the buffer values are carried forward to the new version.

The capacity buffer value that you specify for the Allocation model is considered only if you have enabled allocation model in the policy.

The following tables display the capacity buffer that you can define based on the vCenter Adapter object types:

Object Type	Valid Models for Capacity Buffer
CPU	Demand Allocation
Memory	Demand Allocation
Disk Space	Demand Allocation

Compliance Details

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the compliance for the object types in your policy.

Where You Override the Policy Compliance

To view and override the policy compliance setting, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Compliance**

View the compliance thresholds and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 4-125. Compliance Settings in the Create or Edit Policies Workspace

Option	Description
Lock icon	Enables you to override the policy element settings so that you can customize the policy to monitor the objects in your environment.
Compliance	Allows you to set the compliance score threshold based on the number of violations against those standards.

Workload Automation Details

You can set the workload automation options for your policy, so that vRealize Operations Manager can optimize the workload in your environment as per your definition.

How the Workload Automation Workspace Works

You click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

Where You Set the Policy Workload Automation

To set the workload automation for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Workload Automation**.

Table 4-126. Workload Automation in the Create or Edit Policies Workspace

Option	Description
Workload Optimization	<p>Select a goal for workload optimization.</p> <p>Select Balance when workload performance is your first goal. This approach proactively moves workloads so that the resource utilization is balanced, leading to maximum headroom for all resources.</p> <p>Select Moderate when you want to minimize the workload contention.</p> <p>Select Consolidate to proactively minimize the number of clusters used by workloads. You might be able to repurpose resources that are freed up. This approach is good for cost optimization, while making sure that performance goals are met. This approach might reduce licensing and power costs.</p>
Cluster Headroom	<p>Headroom establishes a required capacity buffer, for example, 20 percent. It provides you with an extra level of control and ensures that you have extra space for growth inside the cluster when required. Defining a large headroom setting limits the systems opportunities for optimization.</p> <p>Note vSphere HA overhead is already included in useable capacity and this setting does not impact the HA overhead.</p>
Advanced Settings	<p>Click Advanced Settings to select what type of virtual machines vRealize Operations Manager moves first to address workload. You can set Storage vMotion on or off. The default is ON.</p>

Metrics and Properties Details

You can select the attribute type to include in your policy so that vRealize Operations Manager can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You enable or disable each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

How the Collect Metrics and Properties Workspace Works

When you create or customize a policy, you can override the base policy settings to have vRealize Operations Manager collect the data that you intend to use to generate alerts, and report the results in the dashboards.






To define the metric and super metric symptoms, metric event symptoms, and property symptoms, in the menu, click **Alerts** and then in the left pane click **Configuration > Symptom Definitions**.

Where You Override the Policy Attributes

To override the attributes and properties settings for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

You can also edit the metrics and properties while working on the objects under the Environment Tab. In the **Metrics** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Metrics Collection**.

Table 4-127. Metrics and Properties Options

Option	Description
Actions	Select one or more attributes and select enable, disable, or inherit to change the state and KPI for this policy.
Filter options	<p>Deselect the options in the Attribute Type, State, KPI, and DT drop-down menus, to narrow the list of attributes.</p> <ul style="list-style-type: none"> ■  Enabled. Indicates that an attribute will be calculated. ■  Enabled (Force). Indicates state change due to a dependency. ■  Disabled. Indicates that an attribute will not be calculated. ■  Inherited. Indicates that the state of this attribute is inherited from the base policy and will be calculated. ■  Inherited. Indicates that the state of this attribute is inherited from the base policy and will not be calculated. <p>The KPI determines whether the metric, property, or super metric attribute is considered to be a key performance indicator (KPI) when vRealize Operations Manager reports the collected data in the dashboards. Filter the KPI states to display attributes with KPI enabled, disabled, or inherited for the policy.</p>
Object Type	Filters the attributes list by object type.
Page Size	The number of attributes to list per page.
Attributes data grid	<p>Display the attributes for a specific object type.</p> <ul style="list-style-type: none"> ■ Name. Identifies the name of the metric or property for the selected object type. ■ Type. Distinguishes the type of attribute to be either a metric, property, or super metric. ■ Adapter Type. Identifies the adapter used based on the object type selected, such as Storage Devices. ■ Object Type. Identifies the type of object in your environment, such as StorageArray. ■ State. Indicates whether the metric, property, or super metric is inherited from the base policy. ■ KPI. Indicates whether the key performance indicator is inherited from the base policy. If a violation against a KPI occurs, vRealize Operations Manager generates an alert. ■ DT. Indicates whether the dynamic threshold (DT) is inherited from the base policy.

Alert and Symptom Details

You can enable or disable alert and symptom definitions to have vRealize Operations Manager identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

How the Alert and Symptom Definitions Workspace Works

vRealize Operations Manager collects data for objects and compares the collected data to the alert definitions and symptom definitions defined for that object type. Alert definitions include associated symptom definitions, which identify conditions on attributes, properties, metrics, and events.

You can configure your local policy to inherit alert definitions from the base policies that you select, or you can override the alert definitions and symptom definitions for your local policy.

Before you add or override the alert definitions and symptom definitions for a policy, familiarize yourself on the available alerts and symptoms.

- To view the available alert definitions, in the menu, click **Alerts** and then in the left pane click **Configuration > Alert Definitions**.
- To view the available symptom definitions, in the menu, click **Alerts** and then in the left pane click **Configuration > Symptom Definitions**. Symptom definitions are available for metrics, properties, messages, faults, smart early warnings, and external events.

A summary of the number of problem and symptoms that are enabled and disabled, and the difference in changes of the problem and symptoms as compared to the base policy, appear in the Analysis Settings pane of the policies workspace.

Where You Override the Alert Definitions and Symptom Definitions

To override the alert definitions and symptom definitions for your policy, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The definitions appear in the workspace.

You can also edit the alert settings while working on the objects under the Environment Tab. In the **Alerts** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Alerts State**.

Policy Alert Definitions and Symptom Definitions

You can override the alert definitions and symptom definitions for each policy.

■ Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

■ Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can enable or disable the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

How the Policy Alert Definitions Work

vRealize Operations Manager uses problems to trigger alerts. A problem manifests when a set of symptoms exists for an object, and requires you to take action on the problem. Alerts indicate problems in your environment. vRealize Operations Manager generates alerts when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert occurs, vRealize Operations Manager presents the triggering symptoms for you to take action.

Some of the alert definitions include predefined symptoms. When you include symptoms in an alert definition, and enable the alert, an alert is generated when the symptoms are true.

The Alert Definitions pane displays the name of the alert, the number of symptoms defined, the adapter, object types such as host or cluster, and whether the alert is enabled as indicated by **Local**, disabled as indicated by **not Local**, or inherited. Alerts are inherited with a green checkmark by default, which means that they are enabled.

You can automate an alert definition in a policy when the highest priority recommendation for the alert has an associated action.

To view a specific set of alerts, you can select the badge type, criticality type, and the state of the alert to filter the view. For example, you can set the policy to send fault alerts for virtual machines.

Where You Modify the Policy Alert Definitions

To modify the alerts associated with policies, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-128. Alert Definitions in the Create or Edit Policies Workspace

Option	Description
Actions	Select one or more alert definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	<p>Deselect the options in the Type and State drop-down menus, to narrow the list of symptom definitions.</p> <p>Impact indicates the health, risk, and efficiency badges to which the alerts apply.</p> <p>Criticality indicates the information, critical, immediate, warning, or automatic criticality types to which the alert definition applies.</p> <p>Automate indicates the actions that are enabled for automation when an alert triggers, or actions that are disabled or inherited. Actions that are enabled for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to Local with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark.</p>
Object Type	Filters the alert definitions list by object type.
Page Size	The number of alert definitions to list per page.
Filter	Locates data in the alert definition list.
Alert Definitions data grid	<p>Displays information about the alert definitions for the object types. The full name for Alert definition and the criticality icon appear in a tooltip when you hover the mouse over the Alert Definition name.</p> <ul style="list-style-type: none"> ■ Alert Definition. Meaningful name for the alert definition. ■ Criticality. Indicates the criticality of the alert. ■ Symptom. Number of symptoms defined for the alert. ■ Actionable Recommendations. Only recommendations with actions in the first priority, as they are the only ones you can automate. ■ Automate. When the action is set to Local, the action is enabled for automation when an alert triggers. Actions that are enabled for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to Local with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ State. Alert definition state, either enabled, disabled, or inherited from the base policy.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can enable or disable the symptom definitions in your policy.

How the Policy Symptom Definitions Work

vRealize Operations Manager uses symptoms that are enabled to generate alerts. When the symptoms used in an alert definition are true, and the alert is enabled, an alert is generated.

When a symptom exists for an object, the problem exists and requires that you take action to solve it. When an alert occurs, vRealize Operations Manager presents the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

To assess objects for symptoms, you can include symptoms packages in your policy for metrics and super metrics, properties, message events, and faults. You can enable or disable the symptoms to determine the criteria that the policy uses to assess and evaluate the data collected from the objects to which the policy applies. You can also override the threshold, criticality, wait cycles, and cancel cycles.

The Symptoms pane displays the name of the symptom, the associated management pack adapter, object type, metric or property type, a definition of the trigger such as for CPU usage, the state of the symptom, and the trigger condition. To view a specific set of symptoms in the package, you can select the adapter type, object type, metric or property type, and the state of the symptom.

When a symptom is required by an alert, the state of the symptom is enabled, but is dimmed so that you cannot modify it. The state of a required symptom includes an information icon that you can hover over to identify the alert that required this symptom.

Where You Modify the Policy Symptom Definitions

To modify the policy package of symptoms, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 4-129. Symptom Definitions in the Create or Edit Policies Workspace






Option	Description
Actions	Select one or more symptom definitions and select enable, disable, or inherit to change the state for this policy.
Filter options	<p>Deselect the options in the Type and State drop-down menus, to narrow the list of symptom definitions.</p> <ul style="list-style-type: none"> ■  Enabled. Indicates that a symptom definition will be included. ■  Enabled (Force). Indicates state change due to a dependency. ■  Disabled. Indicates that a symptom definition not be included. ■  Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will be included. ■  Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will not be included. <p>Type determines whether symptom definitions that apply to HT and DT metrics, properties, events such as message, fault, and metric, and smart early warnings appear in the list.</p> <p>State determines whether enabled, disabled, and inherited symptom definitions appear in the symptom definition list.</p>

Table 4-129. Symptom Definitions in the Create or Edit Policies Workspace (continued)

Option	Description
Object Type	Filters the symptom definitions list by object type
Page Size	The number of symptom definitions to list per page.
Filter	Locate data in the symptom definition list.
Symptom Definitions data grid	<p>Displays information about the symptom definitions for the object types. The full name for Symptom Definition appears in a tooltip when you hover the mouse over the Symptom Definition name.</p> <ul style="list-style-type: none"> ■ Symptom Definition. Symptom definition name as defined in the list of symptom definitions in the Content area. ■ Criticality. Indicates the criticality. ■ Adapter. Data source type for which the alert is defined. ■ Object Type. Type of object to which the alert applies. ■ Type. Object type on which the symptom definition must be evaluated. ■ Trigger. Static or dynamic threshold, based on the number of symptom definitions, the object type and metrics selected, the numeric value assigned to the symptom definition, the criticality of the symptom, and the number of wait and cancel cycles applied to the symptom definition. ■ State. Symptom definition state, either enabled, disabled, or inherited from the base policy. ■ Condition. Enables action on the threshold. When set to Override, you can change the threshold. Otherwise set to default. ■ Threshold. To change the threshold, you must set the State to Enabled, set the condition to Override, and set the new threshold in the Override Symptom Definition Threshold dialog box.

If you do not configure the package, the policy inherits the settings from the selected base policy.

Groups and Objects details

You can assign your local policy to one or more objects or groups of objects to have vRealize Operations Manager analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

How the Groups and Objects Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more objects or groups of objects. vRealize Operations Manager uses the settings in the policy to analyze and collect data from the associated objects, and displays the data in dashboards, views, and reports.

Where You Apply a Policy to Groups and Objects

To apply the policy to an object or groups of objects, in the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**.

Groups and Objects Options

To apply the policy to an object or groups of objects, select the check box for the groups or objects in the workspace.

You can then view the groups and objects associated with the policy. In the menu, click **Administration**, and then in the left pane click **Policies**. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**. Click the **Custom Groups** tab to apply the policy to one or more groups of objects. Click the **Objects** tabs to apply the policy to one or more objects.

For more information about how to create an object group, see the topic called **Custom Object Groups Workspace to Create a New Group**.

For more information about how to create a policy, see [Policy Workspace in vRealize Operations Manager](#).

Define Monitoring Goals for vRealize Operations Manager Solutions

The Manage Solution configuration for the vSphere solution provides a set of questions for you to answer to help you define the default policy settings associated with your vCenter Adapter. You can create a policy for a management pack solution that you add to vRealize Operations Manager.

How Define Monitoring Goals Works in vRealize Operations Manager

The Manage Solution workspace includes an option to define monitoring goals for the solution. The selections you make determine the default policy settings that vRealize Operations Manager uses to analyze and monitor the objects associated with the solution.

For example, you might have a production environment that is composed of four separate production areas, each of which includes specific object groups. To monitor the objects in each production area, you must set the default policy settings according to the monitoring requirements for each area. You can have vRealize Operations Manager set the default settings based on your infrastructure or virtual machines, alert you on individual objects or object groups, and so on.

Where You Define the Monitoring Goals for a Solution

To define the monitoring goals for a solution and establish the default settings for monitoring goals in the default policy, in the menu, click **Administration**, and then in the left pane, click **Solutions > Configuration**, and select a solution. Click **Configure**, and click **Define Monitoring Goals**. In the Define Monitoring Goals dialog box that appears, select answers to the questions about your objects, alerts, memory capacity, and compliance settings according to the *vSphere Hardening Guide*.

When you select an option, vRealize Operations Manager saves your setting. If you display the Define Monitoring Goals dialog box later, and the user interface did not appear to retain your selection, the selection is still active. As a double-check, select the option again, and click **Save**.

To adjust advanced settings of the policy, in the menu, click **Administration**, and then in the left pane, click **Policies**.

Table 4-130. Define Monitoring Goals Questions

Option	Description
Which objects do you want to be alerted on in your environment?	Select the type of objects to receive alerts. You can have vRealize Operations Manager alert on all infrastructure objects except for virtual machines, only virtual machines, or all.
Which type of alerts do you want to enable?	You can enable vRealize Operations Manager to trigger Health, Risk, and Efficiency alerts on your objects.
Configure Memory Capacity based on?	Set the memory capacity model based on the type of environment to monitor. For example, to monitor a production environment, select the vSphere Default model to use moderate settings to ensure performance. Use Most Aggressive for test and development environments. Use Most Conservative to use all allocated memory for capacity calculations.
Enable <i>vSphere Hardening Guide</i> Alerts?	Use the <i>vSphere Hardening Guide</i> to continuously and securely assess and operate your vSphere objects. When you enable these alerts, vRealize Operations Manager assesses your objects against the <i>vSphere Hardening Guide</i> 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.

What Are Compliance Benchmarks

Compliance benchmarks display score cards that help you proactively detect compliance problems in vRealize Operations Manager. The compliance benchmarks are measured against a set of standard rules, regulatory best practices, or custom alert definitions.

How Compliance Benchmarks Work

All the compliance standards in vRealize Operations Manager, including any standards that you define, are based on alert definitions. Only alert definitions of the Compliance subtype are counted. Custom score cards can monitor user-defined alerts.

In previous releases of vRealize Operations Manager, you had to modify the current default policy to monitor compliance against a set of standard rules, regulatory best practices, or custom alert definitions. In the current release, you can manage all compliance related tasks from the **Home > Troubleshoot > Compliance** page. When you configure a benchmark, you select an applicable policy. vRealize Operations Manager then enables the appropriate alert definitions in the policy to measure compliance.

The compliance assessment is based on the environment where your objects are deployed. You can monitor objects that are deployed in your VMware Self-Managed Cloud (SDDC) environment, including DC and Edge environments, and your VMware Managed Cloud (VMC SDDC) environment. Compliance benchmarks on VMC SDDC are applicable only on client VMs that you have deployed in the VMware Managed Cloud environment.

vRealize Operations Manager Compliance Benchmark Types

VMware SDDC Benchmarks

Displays score cards based on alerts which are measured against the latest hardening guides:

- vSphere Security Configuration Guide
- vSAN Security Configuration Guide
- NSX Security Configuration Guide

Displays benchmarks for and in the SDDC and VMC SDDC tabs.

Note vSphere 6.7 Update 1 Security Configuration Guide no longer contains risk profiles. For more information, see blogs.vmware.com.

Custom Benchmarks

Displays benchmarks that you define. Use compliance alerts from vSphere and regulatory management packs, or define your own alerts to monitor. You can define up to five custom score cards. You can import custom score cards from other instances of vRealize Operations Manager.

Regulatory Benchmarks

Displays benchmarks for industry standard regulatory compliance requirements. You can install compliance packs for the following regulatory standards:

- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry Data Security Standard (PCI DSS) compliance standards
- CIS Security Standards
- Defense Information Systems Agency (DISA) Security Standards
- The Federal Information Security Management Act (FISMA) Security Standards
- International Organization for Standardization (ISO) Security Standards

For instructions on installing these compliance packs, see [Install a Regulatory Benchmark](#).

Compliance Score Cards

The compliance page in vRealize Operations Manager displays score cards for each type of benchmark. A score card is a compliance visualization term.

What is a Compliance Score Card

Score cards in the Compliance landing page display the number of non-compliant objects, and the total number of objects affected by each hardening guide as well as the compliance score which is counted as the ratio of compliant objects to total number of objects assessed by the given benchmark, represented in percentage. In addition, you can see the breakdown of the total number of objects that are compliant and non-compliant. You can click on a score card to see more details, including alerts that were triggered based on the compliance standards.

The compliance score card of an object is counted as the smallest rounded off integer ($100 * (\text{total number of symptoms triggered on an object} / \text{total number of symptoms})$).

The compliance score for the object is based on the most critical of the violated standards. The score card displays 100 when all objects are compliant. When an object is non-compliant, the number of non-compliant symptoms are displayed in red and the total number of symptoms in grey.

Where You Find Compliance Score Cards

You can view score cards for each of the different types of benchmarks in the **Home > Troubleshoot > Compliance** page.

You can view score cards for objects in the **Environment > Object > Compliance** tab.

Compliance Page

In the **Home > Troubleshoot > Compliance** summary page, vRealize Operations Manager monitors compliance for SDDC and VMC SDDC objects. You can switch between the tabs to view the benchmarks for your on-premise deployment and cloud environments.

In each of these tabs, vRealize Operations Manager displays compliance score cards in the following sections:

- VMware SDDC Benchmarks
- Custom Benchmarks
- Regulatory Benchmarks

Compliance Tab

In the **Environment > Object > Compliance** tab, vRealize Operations Manager displays score cards for the benchmarks that include the current objects in their calculations, based on the alert definitions and policies associated with that benchmark. The score cards display the total number of rules and the number non-compliant (violated) rules based on symptoms for each hardening guide.

Score Cards in the Compliance Page

In the **Home > Troubleshoot > Compliance** page, you can view scores for benchmarks that you have enabled. Click a score card to view more information.

Table 4-131. Compliance Page Score Card Options

Item	Description
Score card for the configured hardening guides, custom benchmark and management packs	Displays the compliance score, total compliant and non-compliant objects for the compliance standards you have configured.
Object Breakdown	<p>Displays the number of compliant and non-compliant objects for the following types of objects:</p> <ul style="list-style-type: none"> ■ vCenter ■ ESXi Host ■ Virtual Machine ■ Distributed Port Group ■ Distributed Virtual Switch ■ vSAN Cache Disk ■ vSAN Capacity Disk ■ vSAN Cluster ■ NSX-T Manager ■ NSX-V EDGE ■ NSX-V Logical Router ■ NSX-V Manager ■ NSX-V Routing Edge Service
Compliance Alert List	<p>A list of alerts, grouped by time by default. You can either remove the grouping of the alerts, or group by criticality, definition, and object type.</p> <p>The alerts which caused the compliance violation are displayed in a table. You can sort the table by the following columns:</p> <ul style="list-style-type: none"> ■ Alert ID ■ Criticality ■ Alert ■ Triggered On ■ Updated On <p>Select an alert from the table and click Actions to perform tasks such as canceling the alert, suspending alert, and taking ownership of the alert.</p> <p>Click an alert to view more details. The Environment > Object > Alert tab opens.</p>

Compliance Alerts

You use the compliance score card as an investigative tool when you evaluate the state of objects in your environment, or when you research the root cause of a problem. If the score card indicates a problem, you can view the alerts to see details about the violation. Violated rules are based on the symptoms defined in the compliance alert.

The compliance alerts, which have the subtype named Compliance, include one or more symptoms that represent the compliance rules. Compliance alerts that are triggered appear on the **Environment > Object > Compliance** tab as violations to the standard, and the triggered symptoms appear as violated rules. The rules are the alert symptoms, and the symptom configuration identifies the incorrect value or configuration. If a rule symptom is triggered for any of the alerts in the standard, the triggered rule violates the standard and affects the score that appears on the **Environment > Object > Compliance** tab.

Table 4-132. Compliance Tab Alert Display

Item	Description
Score card for the configured hardening guides	Displays the score card value, total number of rules, and number of non-compliance rules for the compliance standards you have configured.
Active Compliance Alerts	<p>If you click the score card, the rules for the score card appear. When a symptom is triggered, the rule is considered to be violated. View the list of rules in the following tabs:</p> <ul style="list-style-type: none"> ■ Violated Rules. Displays only the triggered symptoms. Click a symptom to view more information. ■ All Rules. Displays triggered and untriggered symptoms.

How To Configure Compliance Benchmarks

Configure VMware SDDC, custom, and regulatory benchmarks from the Compliance page. Unlike previous releases, you can now enable alert definitions in one of the active policies, from the Compliance page directly.

Enable VMware SDDC Benchmarks

You can enable the VMware SDDC Benchmark to monitor objects for violation of vSphere Security Configuration Guide, vSAN Security Configuration Guide, NSX Security Configuration Guide (SDDC only). The score cards in the VMware SDDC Benchmark warn you when compliance alerts trigger on your vCenter Server instance, NSX-V objects, NSX-T objects, vSAN objects, ESXi hosts, virtual machines, distributed port groups, or distributed virtual switches.

Procedure

- 1 Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.
- 2 To enable the Security Configuration Guides, select either the SDDC or the VMC SDDC tab depending on the environment where your objects are present.
- 3 In the VMware SDDC Benchmarks section, click **Enable** under the vSphere Security Configuration Guide or vSAN Security Configuration Guide pane.

Note To enable the NSX Security Configuration guide, you must first install the NSX for vSphere, or the NSX-T solution. For more details, see [Adding Solutions](#).

The **Enable Policies** dialog box opens.

- 4 Select the policy that you want to modify. When there are child policies, you can select a child policy and unselect a parent policy. vRealize Operations Manager modifies the selected policy and enables the alert definitions associated with the current scorecard.
- 5 Click **Enable** to confirm your selection.

Results

vRealize Operations Manager starts to assess the objects based on the policy that you selected. To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

Create a New Custom Benchmark

You can create a custom compliance benchmark to ensure that objects comply with compliance alerts available in vRealize Operations Manager, or custom compliance alert definitions. When a compliance alert is triggered on your vCenter instance, hosts, virtual machines, distributed port groups, or distributed switches, you investigate the compliance violation. You can add up to five custom compliance score cards.

Prerequisites

To create a custom benchmark based on industry standard regulatory compliance requirements, you must first download and install the compliance management packs.

Procedure

- 1 Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.
- 2 To create a custom benchmark, first select either the SDDC or the VMC SDDC tab depending on where your objects are present.
- 3 In the Custom Benchmarks section, click **Add Custom Compliance**.
The **Add Custom Compliance** dialog box opens.
- 4 Select **Create a New Custom Benchmark**.
 - a In the Name and Description step, provide a name and description for the custom benchmark and click **Next**.
 - b In the Alert Definitions step, select the compliance alerts that you want to add to this custom compliance benchmark and click **Next**.
 - c In the Policies step, select the policies to enable compliance and click **Finish**.

Results

The custom compliance which monitors alert definitions that you selected is available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit**.

Import or Export a Custom Benchmark

You can export custom benchmarks from any vRealize Operations Manager instance and import it to another instance. Reusing custom benchmarks saves you time and effort. You can modify an imported custom benchmark. Exported files are in the XML format. The XML file contains information about alert groups, alerts, and filters.

Prerequisites

You must first export a XML file with the custom benchmarks from another instance of vRealize Operations Manager before importing the XML file to another instance.

Procedure

- 1 Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance** page.
- 2 To import a custom benchmark, select either the SDDC or the VMC SDDC tab depending on where your objects are present.
- 3 In the Custom Benchmarks section, click **Add Custom Compliance**.
The **Add Custom Compliance** dialog box opens.
- 4 Select **Import An Existing Custom Benchmark**.
 - a In the Import Compliance Score card dialog box, select the Score card Definition XML file from your local computer. If the XML file contains cloned alerts from the vRealize Operations Manager instance that was used to export the file, the cloned alerts are also imported.
 - b vRealize Operations Manager displays a message to indicate if the XML file was successfully imported.
 - c If you see a message which indicates that there is a conflict between the data in the XML file and the custom benchmarks already defined, make a selection on how to handle a conflict.
 - d Click **Done**.
- 5 To export an existing custom benchmark, click the score card to select the benchmark and select **Export** from the **Actions** menu.

Results

The imported compliance benchmarks are available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit** from the **Actions** menu after clicking the score card.

Install a Regulatory Benchmark

To enforce and report on the compliance of your vSphere objects, you activate the compliance pack that contains the policies for regulatory standards. Then, you select the policy to enable the appropriate regulatory alerts for your virtual machines.

Procedure

- 1 Navigate to the Compliance homepage from **Home > Troubleshoot > Compliance**.
The compliance packs for the regulatory standards are displayed under the Regulatory Benchmark section.
- 2 To install any regulatory benchmark, click **Activate From Repository** on the required compliance pack.
You are redirected to the **Native Management Packs** page.
- 3 Navigate to the required compliance pack and click **Activate** to complete the installation.
- 4 To enable the compliance pack policies, navigate to the **Compliance** homepage and click **Enable** on the installed compliance pack.
The **Enable Policies** window opens.
- 5 Select the policies that you want to enable and click **Enable** to complete the process.

Results

vRealize Operations Manager starts to assess the objects based on the regulatory benchmark that you installed.

Configuring Super Metrics

The super metric is a mathematical formula that contains one or more metrics or properties. It is a custom metric that you design to help track combinations of metrics or properties, either from a single object or from multiple objects. If a single metric does not inform you about the behavior of your environment, you can define a super metric.

After you define it, you assign the super metric to one or more object types. This action calculates the super metric for the objects in that object type and simplifies the metrics display. For example, you define a super metric that calculates the average CPU usage on all virtual machines, and you assign it to a cluster. The average CPU usage on all virtual machines in that cluster is reported as a super metric for the cluster.

When the super metric attribute is enabled in a policy, you can also collect super metrics from a group of objects associated with a policy.

Because super metric formulas can be complex, plan your super metric before you build it. The key to creating a super metric that alerts you to the expected behavior of your objects is knowing your own enterprise and data. Use this checklist to help identify the most important aspects of your environment before you begin to configure a super metric.

Table 4-133. Designing a Super Metric Checklist

<input type="checkbox"/> Determine the objects that are involved in the behavior to track.	When you define the metrics to use, you can select either specific objects or object types. For example, you can select the specific objects VM001 and VM002, or you can select the object type virtual machine.
<input type="checkbox"/> Determine the metrics to include in the super metric.	If you are tracking the transfer of packets along a network, use metrics that refer to packets in and packets out. In another common use of super metrics, the metrics might be the average CPU usage or average memory usage of the object type you select.
<input type="checkbox"/> Decide how to combine or compare the metrics.	For example, to find the ratio of packets in to packets out, you must divide the two metrics. If you are tracking CPU usage for an object type, you might want to determine the average use. You might also want to determine what the highest or lowest use is for any object of that type. In more complex scenarios, you might need a formula that uses constants or trigonometric functions.
<input type="checkbox"/> Decide where to assign the super metric.	You define the objects to track in the super metric, then assign the super metric to the object type that contains the objects being tracked. To monitor all the objects in a group, enable the super metric in the policy, and apply the policy to the object group.
<input type="checkbox"/> Determine the policy to which you add the super metric.	After you create the super metric, you add it to a policy. For more information, refer to Policy Workspace in vRealize Operations Manager .

What Else Can You Do with Super Metrics

- To see the super metrics in your environment, generate a system audit report. For more information, refer to [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.

The **Manage Super Metric** wizard opens.

- 3 Enter a meaningful name for the super metric such as **Worst VM CPU Usage (%)** in the **Name** text box.

Note It is important that you have an intuitive name as it appears in dashboards, alerts, and reports. For meaningful names, always use space between words so that it is easier to read. Use title case for consistency with the out of the box metrics and add the unit at the end.

- 4 Provide a brief summary of the super metric in the **Description** text box.

Note Information regarding the super metric, like why it was created and by whom can provide clarity and help you track your super metrics with ease.

- 5 Select the unit of the super metrics from the **Unit** drop-down and click **Next**.

Note The super metrics unit configured here can be changed in the metrics charts, widgets, and views.

The Create a formula screen appears.

- 6 Create the formula for the super metric.

For example, to add a super metric that captures the average CPU usage across all virtual machines in a cluster, perform the following steps.

- a Select the function or operator. This selection helps combine the metric expression with operators and/or functions. In the super metric editor, enter **avg** and select the **avg** function.

You can manually enter functions, operators, objects, object types, metrics, metrics types, property, and properties types in the text box and use the suggestive text to complete your super metric formula.

Alternatively, select the function or operator from the **Functions** and **Operators** drop-down menus.

- b To create a metric expression, enter **Virtual** and select **Virtual Machine** from the object type list.

- c Add the metric type, enter **usage**, and select the **CPU|Usage (%)** metric from the metric type list.

Note The expression ends with depth=1 by default. If the expression ends with depth=1, that means that the metric is assigned to an object that is one level above virtual machines in the relationship chain. However, since this super metric is for a cluster which is two levels above virtual machine in the relationship chain, change the depth to 2.

The depth can also be negative, this happens when you need to aggregate the parents of a child object. For example, when aggregating all the VMs in a datastore, the metric expression ends with depth=-1, because VM is a parent object of datastore. But, if you want to aggregate all the VMs at a Datastore Cluster level, you need to implement 2 super metrics. You cannot directly aggregate from VM to Datastore Cluster, because both are parents of a datastore. For a super metric to be valid, depth cannot be 0 (-1+1=0). Hence, you need to create the first super metric (with depth=-1) for the aggregate at the datastore level, and then build the second super metric based on the first (with depth = 1).

The metric expression is created.

- d To calculate the average CPU usage of powered on virtual machines in a cluster, you can add the `where` clause. Enter **where=""**.

Note The **where** clause cannot point to another object, but can point to a different metric in the same object. For example, you cannot count the number of VMs in a cluster with the CPU contention metric > SLA of that cluster. The phrase "SLA of that cluster " belongs to the cluster object, and not to the VM object. The right operand must also be a number and cannot be another super metric or variable. The where clause cannot be combined using AND, OR, NOT, which means you cannot have `where="VM CPU>4 and VM RAM>16"` in your super metric formula.

- e Position the pointer between the quotation marks, enter **Virtual**, and select the **Virtual Machine** object type and the **System|Powered ON** metric type.
- f To add the numeric value for the metric, enter **==1**.
- g To view hints and suggestions, click **ctrl+space** and select the adapter type, objects, object types, metrics, metrics types, property, and properties types to build your super metric formula.
- h Click the **This object** icon.

If the **This object** icon is selected during the creation of a metric expression, it means that the metric expression is associated to the object for which the super metric is created.

- 7 You can also use the **Legacy** template to create a super metric formula without the suggestive text.

To view the super metric formula in a human-readable format, click the **Show Formula Description** icon. If the formula syntax is wrong, an error message appears.

Note If you are using Internet Explorer, you are automatically directed to the legacy template.

- 8 Verify that the super metric formula has been created correctly.

- a Expand the **Preview** section.

- b In the **Objects** text box, enter and select a **Cluster**.

A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

- c Click the **Snapshots** icon.

You can save a snapshot, or download the metric chart in a `.csv` format.

- d Click the **Monitoring Objects** icon.

If enabled, only the objects that are being monitored are used in the formula calculation.

- e Click **Next**.

The Assign to Object Types screen appears.

- 9 Associate the super metric with an object type. vRealize Operations Manager calculates the super metric for the target objects and displays it as a metric for the object type.

- a In the **Assign to an Object Type** text box, enter **Cluster** and select the **Cluster Compute Resource** object type.

After one collection cycle, the super metric appears on each instance of the specified object type. For example, if you define a super metric to calculate the average CPU usage across all virtual machines and assign it to the cluster object type, the super metric appears as a super metric on each cluster.

- b Click **Next**.

The Enable in a Policy screen appears.

- 10 Enable the super metric in a policy, wait for at least one collection cycle till the super metric begins collecting and processing data, and then review your super metric on the **All Metrics** tab.

- a In the **Enable in a Policy** section, you can view the policies related to the object types you assigned your super metric to. Select the policy in which you want to enable the super metric. For example, select the **Default Policy** for Cluster.

- 11 Click **Finish**.

You can now view the super metric you created and the associated object type and policy on the **Super Metrics** page.

Enhancing Your Super Metrics

You can enhance your super metrics by using clauses and resource entry aliasing.

Where Clause

The **where** clause verifies whether a particular metric value can be used in the super metric. Use this clause to point to a different metric of the same object, such as **where=(\$ {metric=metric_group|my_metric} > 0)**.

For example: `count($ {objecttype = ExampleAdapter, adaptertype = ExampleObject, metric = ExampleGroup|Rating, depth=2, where = ($value==1)})`

IsFresh Function

Use the **isFresh** function in the **where** clause to check if the last value of the metrics is fresh or not.

For every metric published in vRealize Operations Manager, the point with the latest publishing time is called as the last point of that metric. The value of that metric's last point is called the last value of that metric. A metric's last point is considered fresh when the time elapsed after the metric's last point is lesser than the estimated publishing interval of that metric.

The **isFresh** function returns true if the last value of the metrics is fresh. For example, in the following scenarios, the function:

- `${this, metric=a|b, where=($value.isFresh())}`, returns the last value of the metric a|b if the last value is fresh.
- `${this, metric=a|b, where=($value == 7 && $value.isFresh())}`, returns the last value of the metric a|b if it is equal to seven and is fresh.
- `${this, metric=a|b, where=($ {metric=c|d} == 7 && $ {metric=c|d}.isFresh())}`, returns the last value of the metric a|b only if the last value of the metric c|d is equal to seven and is fresh.

Resource Entry Aliasing

Resource entries are used to retrieve metric data from vRealize Operations Manager for computing super metrics. A resource entry is the part of an expression which begins with `$` followed by a `{ . . }` block. When computing a super metric, you might have to use the same resource entry multiple times. If you have to change your computation, you must change every resource entry, which might lead to errors. You can use resource entry aliasing to rewrite the expression.

The following example, shows a resource entry that has been used twice.

```
(min($ {adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)/
(max($ {adaptertype=VMWARE, objecttype=HostSystem, attribute=cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)"
```

The following example shows how to write the expressing using resource entry aliasing. The output of both expressions is the same.

```
(min(${adaptype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)} as cpuload) + 0.0001)/
(max(cpuload) + 0.0001) "
```

Follow these guidelines when you use resource entry aliasing:

- When you create an alias, make sure that after the resource entry you write **as** and then **alias:name**. For example: **\${...} as alias_name**.
- The alias cannot contain the ()[]+~*/%|&!=<>.,?:\$ special characters, and cannot begin with a digit.
- An alias name, like all names in super metric expressions, is case-insensitive.
- Use of an alias name is optional. You can define the alias, and not use it in an expression.
- Each alias name can be used only once. For example: **\${resource1,...} as r1 + \${resource2,...} as R1**.
- You can specify multiple aliases for the same resource entry. For example: **\${...} as a1 as a2**.

Conditional Expression ? : Ternary Operators

You can use a ternary operator in an expression to run conditional expressions.

For example: **expression_condition ? expression_if_true : expression_if_false**.

The result of the conditional expression is converted to a number. If the value is not 0, then the condition is assumed as true.

For example: **-0.7 ? 10 : 20** equals 10. **2 + 2 / 2 - 3 ? 4 + 5 / 6 : 7 + 8** equals 15 (7 + 8).

Depending on the condition, either **expression_if_true** or **expression_if_false** is run, but not both of them. In this way, you can write expressions such as, **\${this, metric=cpu|demandmhz} as a != 0 ? 1/a : -1**. A ternary operator can contain other operators in all its expressions, including other ternary operators.

For example: **!1 ? 2 ? 3 : 4 : 5** equals 5.

Exporting and Importing a Super Metric

You can export a super metric from one vRealize Operations Manager instance and import it to another vRealize Operations Manager instance. For example, after developing a super metric in a test environment, you can export it from the test environment and import it use in a production environment.

If the super metric to import contains a reference to an object that does not exist in the target instance, the import fails. vRealize Operations Manager returns a brief error message and writes detailed information to the log file.

Procedure

1 Export a super metric.

- a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.
- b Select the super metric to export, click the **Actions** icon and select **Export Selected Super Metric** icon.

vRealize Operations Manager creates a super metric file, for example, `SuperMetric.json`.
- c Download the super metric file to your computer.

2 Import a super metric.

- a On the menu, select **Administration** and in the left pane select **Configuration > Super Metrics**.
- b Click the **Actions** icon and select **Import Super Metric**.
- c (Optional). If the target instance has a super metric with the same name as the super metric you are importing, you can either overwrite the existing super metric or skip the import, which is the default.

Super Metrics Tab

A super metric is a mathematical formula that contains a combination of one or more metrics for one or more objects. With super metrics you can assess information more quickly when you are observing fewer metrics.

Where You Configure Super Metrics

Click **Administration** and in the left pane click **Configuration > Super Metrics**.

Table 4-134. Configuration Options for Super Metrics

Option	Description
Toolbar	<p>Use the toolbar selections to manage super metric options.</p> <ul style="list-style-type: none"> ■ Add New Super Metric. Starts the Manage Super Metric workspace. See Manage Super Metric Workspace. ■ Edit Selected Super Metric. Starts the Manage Super Metric workspace. ■ Clone Selected Super Metric. Duplicates the super metric. Edit the clone or associate it with a different object type. ■ Delete Selected Super Metric. ■ Export Selected Super Metric. Exports a super metric to use in another vRealize Operations Manager instance. See Exporting and Importing a Super Metric. ■ Import Super Metric. Imports a super metric to this vRealize Operations Manager instance. See Exporting and Importing a Super Metric.
Super Metrics list	Configured super metrics listed by name and formula description.
Policies Tab	Policies in which the super metric attribute is enabled for collection. When enabled in a policy, vRealize Operations Manager collects super metrics from the objects associated with the policy. See Metrics and Properties Details .
Object Types Tab	Object types for the super metric display. vRealize Operations Manager calculates the super metric for the objects associated with the object type and displays the value with the object type. Use the toolbar selections to add or delete an object type association.

Manage Super Metric Workspace

You use the Manage Super Metric workspace to create or edit a super metric. The toolbar helps you to build the mathematical formula with the objects and metrics you select.

Where You Configure Super Metrics

On the menu, click **Administration** and in the left pane click **Configuration > Super Metrics**.

Table 4-135. Super Metrics Workspace Options

Option	Description
Super Metric	<p>Use the toolbar selections to build and display your super metric formula.</p> <ul style="list-style-type: none"> ■ 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 <code>this</code> 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	<p>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.</p> <ul style="list-style-type: none"> ■ 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 4-136. 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 4-137. Looping Function Formats

Argument syntax example	Description
<code>func({this, metric = a/b:optional_instance/c})</code>	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.
<code>func({this, attribute = a/b:optional_instance/c})</code>	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.
<code>func({adaptype = adapkind, objecttype = reskind, resourcename = resname, identifiers = {id1 = val1, id2 = val2, ...}, metric = a/b:instance/c})</code>	Returns a single data point of a particular metric for the <i>resname</i> specified in the argument. This super metric does not take values from the children or parents of the object.
<code>func({adaptype = adapkind, objecttype = reskind, resourcename = resname, identifiers = {id1 = val1, id2 = val2, ...}, attribute = a/b:optional_instance/c})</code>	Returns a set of data points. This function iterates attributes of the <i>resname</i> specified in the argument. This super metric does not take values from the child or parent of the object.

Table 4-137. Looping Function Formats (continued)

Argument syntax example	Description
<i>func</i> (\$ {adaptype= <i>adaptkind</i> , objecttype= <i>reskind</i> , depth= <i>dep</i> }, metric= <i>a</i> / <i>b:optional_instance/c</i>)	Returns a set of data points. This function iterates metrics of the <i>reskind</i> specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects, where <i>depth</i> 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.
<i>func</i> (\$ {adaptype= <i>adaptkind</i> , objecttype= <i>reskind</i> , depth= <i>dep</i> }, attribute= <i>a</i> / <i>b:optional_instance/c</i>)	Returns a set of data points. This function iterates attributes of the <i>reskind</i> specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects.

For example, `avg ($ {adaptype=VMWARE, objecttype=VirtualMachine, attribute=cpu | usage_average, depth=1})` averages the value of all metric instances with the `cpu | usage_average` attribute for all objects of type `VirtualMachine` that the vCenter adapter finds. vRealize Operations Manager searches for objects one level below the object type where you assign the super metric.

Single Functions

Single functions work on only a single value or a single pair of values.

Table 4-138. Single Functions

Function	Format	Description
<i>abs</i>	<i>abs</i> (x)	Absolute value of x. x can be any floating point number.
<i>acos</i>	<i>acos</i> (x)	Arccosine of x.
<i>asin</i>	<i>asin</i> (x)	Arcsine of x.
<i>atan</i>	<i>atan</i> (x)	Arctangent of x.
<i>ceil</i>	<i>ceil</i> (x)	The smallest integer that is greater than or equal to x.
<i>cos</i>	<i>cos</i> (x)	Cosine of x.
<i>cosh</i>	<i>cosh</i> (x)	Hyperbolic cosine of x.
<i>exp</i>	<i>exp</i> (x)	e raised to the power of x.
<i>floor</i>	<i>floor</i> (x)	The largest integer that is less than or equal to x.
<i>log</i>	<i>log</i> (x)	Natural logarithm (base x) of x.
<i>log10</i>	<i>log10</i> (x)	Common logarithm (base 10) of x.
<i>pow</i>	<i>pow</i> (x,y)	Raises x to the y power.

Table 4-138. Single Functions (continued)

Function	Format	Description
<i>rand</i>	rand()	Generates a pseudo random floating number greater than or equal to 0.0 and less than 1.0.
<i>sin</i>	sin(x)	Sine of x.
<i>sinh</i>	sinh(x)	Hyperbolic sine of x.
<i>sqrt</i>	sqrt(x)	Square root of x.
<i>tan</i>	tan(x)	Tangent of x.
<i>tanh</i>	tanh(x)	Hyperbolic tangent of x.

Operators

Operators are mathematical symbols and text to enclose or insert between functions.

Table 4-139. Numeric Operators

Operators	Description
+	Plus
-	Subtract
*	Multiply
/	Divide
%	Modulo
==	Equal
!=	Not equal
<	Less than
<=	Less than, or equal
>	Greater than
>=	Greater than, or equal
	Or
&&	And
!	Not
? :	<p>Ternary operator. If/then/else</p> <p>For example: conditional_expression ? expression_if_condition_is_true : expression_if_condition_is_false</p> <p>For more information about ternary operators, see Enhancing Your Super Metrics.</p>

Table 4-139. Numeric Operators (continued)

Operators	Description
()	Parentheses
[]	Use in an array of expressions
[x, y, z]	An array containing x, y, z. For example, min([x, y, z])

Table 4-140. String Operators

String Operators	Description
equals	Returns true if metric/property string value is equal to specified string.
contains	Returns true if metric/property string value contains specified string.
startsWith	Returns true if metric/property string value starts with the specified prefix.
endsWith	Returns true if metric/property string value ends with the specified suffix.
!equals	Returns true if metric/property string value is not equal to specified string.
!contains	Returns true if metric/property string value does not contain specified string.
!startsWith	Returns true if metric/property string value does not start with the specified prefix.
!endsWith	Returns true if metric/property string value does not end with the specified suffix.

Note String operators are valid in 'where' condition only. For example: `${this, metric=summary|runtime|isIdle, where = "System Properties|resource_kind_type !contains GENERAL"}`

Configuring Objects

Using the power of object management - including metrics and alerts - you can monitor objects, applications, and systems that must stay up and running. Some metrics and alerts are prepackaged into dashboards and policies; others you combine into custom tools

vRealize Operations Manager discovers objects in your environment and makes them available to you. With the information that vRealize Operations Manager provides, you can quickly access and configure any object. For example, you can determine if a datastore is connected or providing data, or you can power on a virtual machine.

Object Discovery

Its ability to monitor and collect data on objects in your systems environment makes vRealize Operations Manager a critical tool in maintaining system uptime and ensuring ongoing good health for all system resources from virtual machines to applications to storage - across physical, virtual, and cloud infrastructures.

Following are examples of objects that can be monitored.


- vCenter Server
- Virtual machines
- Servers/hosts
- Compute resources
- Resource pools
- Data centers
- Storage components
- Switches
- Port groups
- Datastores

Adapters – Key to Object Discovery

vRealize Operations Manager collects data and metrics from objects using adapters, that are the central components of management packs. You can customize adapter instances for your virtual environment using cloud accounts and other accounts. vRealize Operations Manager uses cloud accounts to manage the communication and integration with other products, applications, and functions.

- Cloud Accounts - You can configure cloud adapter instances and collect data from cloud solutions that are already installed in your cloud environment from the cloud accounts page.
- Other Accounts - You can view and configure native management packs and other solutions that are already installed and configure adapter instances from the other accounts page.
- Repository - You can activate or deactivate native management packs and add or upgrade other management packs from the Repository page.

The screenshot displays the list of available solutions in vRealize Operations Manager . You must first Activate the solution before adding and configuring the accounts.



For complete information on configuring management packs and adapters, see [Connecting vRealize Operations Manager to Data Sources](#)

When you create a new adapter instance, it begins discovering and collecting data from the objects designated by the adapter, and notes the relationships between them. Now you can begin to manage your objects.

Workload Management Inventory Objects

vRealize Operations Manager discovers the following workload management objects and their child objects using the vCenter adapter:

- Tanzu Kubernetes cluster
- vSphere Pods
- Namespace

A cluster with Kubernetes enabled, running on vSphere, is called a Supervisor Cluster. In the vRealize Operations Manager inventory, the summary tab of the Supervisor Cluster indicates that it has workload management enabled. The Supervisor Cluster contains specific objects that enable the capability to run Kubernetes workloads within ESXi. vRealize Operations Manager collects metrics and data for the Supervisor Cluster. Supervisor Clusters contain Namespaces, which are resource pools that have dedicated memory, CPU, and storage.

Namespaces contain virtual machines with k8s enabled. They are called k8s control VMs. These VMs are managed by vSphere. Therefore, you cannot take action on these VMs from within vRealize Operations Manager.

DevOps engineers can run workloads on containers running inside vSphere Pods. They can create Tanzu k8s cluster inside a Namespace. A vSphere Pod is a VM with a small footprint that runs one or more Linux containers. It is the equivalent of a k8s pod. A Tanzu Kubernetes cluster is a full distribution of the open-source [Kubernetes](#) container orchestration software that is packaged, signed, and supported by VMware.

To understand the vSphere Tanzu Kubernetes architecture, see *Configuring and Managing vSphere with Kubernetes* in the vSphere documentation.

Workload management objects are excluded from the following workflows:

- Compliance
- Reclaim
- Rightsizing
- Workload optimization

About Objects

Objects are the structural components of your mission-critical IT applications: virtual machines, datastores, virtual switches and port groups are examples of objects.

Because downtime equals cost - in unused resources and lost business opportunities - it's crucial that you successfully identify, monitor and track objects in your environment. The goal is to proactively isolate, troubleshoot and correct problems even before users are aware that anything is wrong.

When a user actually reports an issue, the solution should be quick and comprehensive.

For a complete list of objects that can be defined in vRealize Operations Manager refer to [Object Discovery](#).

vRealize Operations Manager gives you visibility into objects including applications, storage and networks across physical, virtual and cloud infrastructures through a single interface that relates performance information to positive or negative events in the environment.

Managing Objects

When you monitor a large infrastructure, the number of objects and corresponding metrics in vRealize Operations Manager grows rapidly, especially as you add solutions that extend dynamic monitoring and alerts to more parts of your infrastructure. vRealize Operations Manager gives you ample tools to stay abreast of events and issues.

Adding Objects and Configuring Object Relationships

vRealize Operations Manager automatically discovers objects and their relationships once you create an adapter instance. You have the added ability to manually add any objects that you want monitored and to configure object relationships using abstract concepts rather than the connections recorded by vRealize Operations Manager. Where vRealize Operations Manager might discover the classic parent-child relationships between objects, you can create relationships between objects that might not normally be related. For example, you could configure all the datastores supporting a company department to be related.

When objects are related, a problem with one object appears as an anomaly on related objects. So object relationships can help you to identify problems in your environment quickly. The object relationships that you create are called custom groups.

Custom Groups

To create an automated management system you need some way to organize objects so that you can quickly gain insights. You can achieve a high level of automation using custom groups. You have multiple options for tailoring group attributes to support your monitoring strategy.

For example, you can designate a group either to be static or to be updated automatically with membership criteria that you designate. Consider a non-static group of all virtual machines that are powered on and have OS type Linux. When you power on a new Linux VM, it is automatically added to the group and the policy is applied.

For additional flexibility, you can also specify individual objects to be always included or excluded from a given custom group. Or you can have a different set of alerts and capacity calculations for your production environment versus your testing environments.

Managing Applications

vRealize Operations Manager allows you to create containers or objects that can contain a group of virtual machines or other objects in different structural tiers. This new application can then be managed as a single object, and have health badges and alarms aggregated from the child objects of the group.

For example, the system administrator of an online training system might request that you monitor components in the Web, application and database tiers of the training environment. You build an application that groups related training objects together in each tier. If a problem occurs with one of the objects, it is highlighted in the application display and you can investigate the source of the problem.

The Power of Object Management

Using the power of object management, including metrics and alerts - some prepackaged into dashboards and policies, others that you combine into custom monitoring tools - you'll keep a close watch on the objects, applications and systems that must stay up and running.

Managing Objects in Your Environment

An object is the individual managed item in your environment for which vRealize Operations Manager collects data, such as a router, switch, database, virtual machine, host, and vCenter Server instances.

The system requires specific information about each object. When you configure an adapter instance, vRealize Operations Manager performs object discovery to start collecting data from the objects with which the adapter communicates.

An object can be a single entity, such as a database, or a container that holds other objects. For example, if you have multiple Web servers, you can define a single object for each Web server and define a separate container object to hold all of the Web server objects. Groups and applications are types of containers.

Categorize your objects using tags, so that you can easily find, group, or filter them later. A tag type can have multiple tag values. You or vRealize Operations Manager assigns objects to tag values. When you select a tag value, vRealize Operations Manager displays the objects associated with that tag. For example, if a tag type is Lifecycle and tag values are Development, Test, Pre-production, and Production, you might assign virtual machine objects VM1, VM2, or VM3 in your environment to one or more of these tag values, depending on the virtual machine function.

Adding an Object to Your Environment

You might want to add an object by providing its information to vRealize Operations Manager . For example, some solutions cannot discover all the objects that might be monitored. For these solutions, you must either use manual discovery or manually add the object.

When you add an individual object, you provide specific information about it, including the kind of adapter to use to make the connection and the connection method. For example, a vSAN adapter does not know the location of the vSAN devices that you want to monitor.

Prerequisites

Verify that an adapter is present for the object you plan to add. See [Connecting vRealize Operations Manager to Data Sources](#).

Verify that an adapter is present for the object you plan to add. See the *vRealize Operations Manager vApp Deployment and Configuration Guide* .

Note Objects added to vRealize Operations Manager via API will require an OSI license per object.

Procedure

- 1 In the menu, click **Administration**, then select **Configuration > Inventory** from the left pane.
- 2 On the toolbar, click the plus sign.

- 3 Use the topic menus to reveal all fields and provide the required information.

Option	Description
Display name	Enter a name for the object. For example, enter vSAN-Host1 .
Description	Enter any description. For example, enter vSAN-Host monitored with vSAN adapter .
Adapter type	Select an adapter type. For example, select vSAN Adapter .
Adapter instance	Select an adapter instance.
Object type	Select an object type. For a vSAN adapter, you might select vSAN-Host. When you select the object type, the dialog box selections change to include information you provide so that vRealize Operations Manager can find and connect with the selected object type.
Host IP address	Enter the host IP. For example, enter the IP address of vSAN-Host1.
Port number	Accept the default port number or enter a new value.
Credential	Select the Credential, or click the plus sign to add new login credentials for the object.
Collection interval	Enter the collection interval, in minutes. For example, if you expect the host to generate performance data every 5 minutes, set the collection interval to 5 minutes.
Dynamic Thresholding.	Accept the default, Yes.

- 4 Click **OK** to add the object.

Results

vSAN-Host1 appears in the Inventory as a host object type for the vSAN adapter type.

What to do next

When you add an individual object, vRealize Operations Manager does not begin collecting metrics for the object until you turn on data collection. See [Inventory : List of Objects](#).

For each new object, vRealize Operations Manager assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags. See [Creating and Assigning Tags](#).

For each new object, vRealize Operations Manager assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags.

Configuring Object Relationships

vRealize Operations Manager shows the relationship between objects in your environment. Most relationships are automatically formed when the objects are discovered by an installed adapter. In addition, you can use vRealize Operations Manager to create relationships between objects that might not normally be related.

Objects are related physically, logically, or structurally.

- Physical relationships represent how objects connect in the physical world. For example, virtual machines running on a host are physically related.

- Logical relationships represent business silos. For example, all the storage objects in an environment are related to one another.
- Structural relationships represent a business value. For example, all the virtual machines that support a database are structurally related.

Solutions use adapters to monitor the objects in your environment so that physical relationship changes are reflected in vRealize Operations Manager . To maintain logical or structural relationships, you can use vRealize Operations Manager to define the object relationships. When objects are related, a problem with one object appears as an influence on related objects. So object relationships can help you to identify problems in your environment quickly.

Apart from the parent-child relationship, you can also define new relationships in vRealize Operations Manager . The relationship between objects in your environment can be one-to-many, many-to-one, or one-one, the relationship can be defined in horizontal , vertical, or diagonal levels.

Adding an Object Relationship

Parent-child relationships normally occur between interrelated objects in your environment. For example, a data center object for a vCenter Adapter instance might have datastore, cluster, and host system child objects.

The most common object relationships gather similar objects into groups. When you define a custom group with parent objects, a summary of that group shows alerts for that object and for any of its descendants. You can create relationships between objects that might not normally be related. For example, you might define a child object for an object in the group. You define these types of relationships by configuring object relationships.

Procedure

- 1 At the Home page, select **Administration**. Then select **Configuration > Object Relationships** in the left pane.
- 2 In the Parent Selection column, expand the object tag and select a tag value that contains the object to act as the parent object.

The objects for the tag value appear in the top pane of the second column.

- 3 Select a parent object.

Current child objects appear in the bottom pane of the second column.

- 4 In the column to the right of the List column, expand the object tag and select a tag value that contains the child object to relate to the parent.

- 5 (Optional) If the list of objects is long, filter the list to find the child object or objects.

Option	Action
Navigate the object tag list for an object	Expand the object tag in the pane to the right of the List column and select a tag value that contains the object. The objects for the tag value appear in the List column. If you select more than one value for the same tag, the list contains objects that have either value. If you select values for two or more different tags, the list includes only objects that have all of the selected values.
Search for an object by name	If you know all or part of the object name, enter it in the Search text box and press Enter.

- 6 To make an object a child object of the parent object, select the object from the list and drag it to the parent object in the top pane of the second column, or click the **Add All Objects To Parent** icon to make all of the listed objects children of the parent object.

You can use Ctrl+click to select multiple objects or Shift+click to select a range of objects.

Example: Custom Group with Child Objects

If you want vRealize Operations Manager to monitor objects in your environment to ensure that service level capacity requirements for your IT department are met, you add the objects to a custom group, apply a group policy, and define criteria that affect the membership of objects in the group. If you want to monitor the capacity of an object that does not affect the service level requirements, you can add the object as a child of a parent object in the group. If a capacity problem exists for the child object, the summary of the group shows an alert for the parent object.

Object Relationships Workspace

Objects in an enterprise environment are related to other objects in that environment. Objects are either part of a larger object, or they contain smaller component objects, or both.

How Object Relationships Works

When you select a parent object, vRealize Operations Manager shows any related child objects. You can delete a child object or add more child objects from the list of objects in your environment.

Where You Find Object Relationships

At the Home page, select **Administration**. Then select **Configuration > Object Relationships** in the left pane.

Object Relationships Workspace Options

- Two columns in the center pane display the existing parent-child relationships. You use the object tag options above the left column to select a parent object.
- Two columns in the right pane list objects in your environment. You use the object tag options above the right column to select the object to add as a child.

Table 4-141. Object Tag Options

Option	Description
Collapse all.	Closes all the tag group selections.
Deselect All.	Tags remain selected until deselected. Use this option to deselect all tags.

When a parent object has children, the parent selection shows the child objects and the child object options are active.

Table 4-142. Child Object Options

Option	Description
Clear Selections.	Clear all child object selections.
Select All.	Select all child objects. To remove most child objects from the relationship, use this option then click the child objects you do not want to delete.
Remove Selected Children from Relationship.	Removes the selected children from the relationship.
Remove All Children from Relationship.	Select all children listed on the page and remove them from the relationship.
Per Page.	Number of children to list per page.
Search.	Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.

Use the list options to manage the objects to add as children.

Table 4-143. List Options

Option	Description
Clear Selections.	Clear all object selections.
Select All.	Select all objects displayed.
Add All Objects to Parent.	Select all children listed on the page and add them to the parent.
Per page.	Number of objects to list per page.
Search.	Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.

Creating and Assigning Tags

A large enterprise can have thousands of objects defined in vRealize Operations Manager . Creating object tags and tag values makes it easier to find objects and metrics. With object tags, you select the tag value assigned to an object and view the list of objects that are associated with that tag value.

A tag is a type of information, for example, Adapter Types. Adapter Types is a predefined tag. Tag values are individual instances of that type of information. For example, when the system discovers objects using the vCenter Adapter, it assigns all the objects to the vCenter Adapter tag value under the Adapter Types tag.

You can assign any number of objects to each tag value, and you can assign a single object to tag values under any number of tags. You typically look for an object by looking under its adapter type, its object type, and possibly other tags.

If an object tag is locked, you cannot add objects to it. vRealize Operations Manager maintains locked object tags.

- **Predefined Object Tags**

vRealize Operations Manager includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

- **Add an Object Tag and Assign Objects to the Tag**

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

- **Use a Tag to Find an Object**

The quickest way to find an object in vRealize Operations Manager is to use tags. Using tags is more efficient than searching through the entire object list.

Predefined Object Tags

vRealize Operations Manager includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

For example, when you add an object, the system assigns it to the tag value for the collector it uses and the kind of object that it is. vRealize Operations Manager creates tag values if they do not already exist.

If a predefined tag has no values, there is no object of that tag type. For example, if no applications are defined, the applications tag has no tag values.

Each tag value appears with the number of objects that have that tag. Tag values that have no objects appear with the value zero. You cannot delete the predefined tags or tag values.

Table 4-144. 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 > Management > Licensing. Objects are assigned to the license groups during vRealize Operations Manager installation.
Untag	Drag an object to this tag to delete the tag assignment.

Add an Object Tag and Assign Objects to the Tag

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

Prerequisites

Become familiar with the predefined object tags.

Procedure

- 1 Click **Administration** in the menu, then click **Configuration > Inventory** in the left pane.
- 2 Click the **Manage Tags** icon above the list of tags.
- 3 Click the **Add New Tag** icon to add a new row and type the name of the tag in the row.
For example, type **Cloud Objects** and click **Update**.
- 4 With the new tag selected, click the **Add New Tag Value** icon to add a new row and type the name of the value in the row.
For example, type **Video Cloud** and click **Update**.
- 5 Click **OK** to add the tag.
- 6 Click the tag to which you want to add objects to display the list of object tag values.
For example, click **Cloud Objects** to display the Video Cloud object tag value.
- 7 Drag objects from the list in the right pane of the Inventory onto the tag value name.
You can press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.
For example, if you want to assign datacenters that are connected through the vCenter Adapter, type **vCenter** in the search filter and select the datacenter objects to add.

Use a Tag to Find an Object

The quickest way to find an object in vRealize Operations Manager is to use tags. Using tags is more efficient than searching through the entire object list.

Tag values that can also be tags are Applications and Object Types. For example, the Object Types tag has values for each object that is in vRealize Operations Manager, such as Virtual Machine, which includes all the virtual machine objects in your environment. Each of these virtual machines is also a tag value for the Virtual Machine tag. You can expand the tag value list to select the value for which you want to see objects.

Procedure

- 1 In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane.
- 2 In the tag list in the center pane, click a tag for an object with an assigned value.
When you click a tag, the list of values expands under the tag. The number of objects that is associated with each value appears next to the tag value.
A plus sign next to a tag value indicates that the value is also a tag and that it contains other tag values. You can click the plus sign to see the subvalues.

3 Select the tag value.

The objects that have that tag value appear in the pane on the right. If you select multiple tag values, the objects in the list depend on the values that you select.

Tag Value Selection	Objects Displayed
More than one value for the same tag	The list includes objects that have either value. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, the list shows objects that have either value.
Values for two or more different tags	The list includes only objects that have all of the selected values. For example, if you select two values of the Object Types tag, such as Datacenter and Host System, and you also select an adapter instance such as vC-1 of the vCenter Adapter instance tag, only Datacenter or Host System objects associated with vC-1 appear in the list. Datacenter or Host System objects associated with other adapter instances do not appear in the list, nor do objects that are not Datacenter or Host System objects.

4 Select the object from the list.

Manage Object Tags Workspace

A large enterprise can have thousands of objects. When objects are assigned to a tag, and you choose to display objects with that tag value, the objects are easier to find on the Inventory list.

Where You Find Manage Object Tags

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane.

Click the **Manage Tags** icon above the list of tags in the middle pane.

Manage Object Tags Options

The Manage Object Tags screen appears with previously created tags listed. In the left pane, you add tags. In the right pane, you add tag values.

- Click **Add a New Tag** and type a new tag name, or select a tag to delete.
- For the selected tag, click **Add a New Tag Value** and type a new tag value name, or select a tag value to delete.
- For the GEO Location tag, tag values are identified with a location on a world map. Select the tag value and click **Manage Location** to display the **Manage Location** map and pick a geographical location. Objects assigned to that tag value appear in that geographical location on the [Inventory : Geographical Map of Objects](#).

Manage Object Type Tags Workspace

Every object in your environment is of a particular object type. You use Manage Object Type Tags to control the object type tags displayed.

How Manage Object Type Tags Works

For every adapter instance installed, vRealize Operations Manager discovers objects in your environment and starts collecting data from those objects.

Where You Find Manage Object Type Tags

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane. Click the **Manage Object Type Tags** icon above the list of tags.

Manage Object Type Tags Options

Depending on the number of adapters installed , there may be hundreds of object type tags. The Manage Object Type Tags options allow you to turn on or off the tags listed.

- Type a filter word to show the object type tags with the word.
- Name lists all the object type tags.
- To toggle the display of an object type tag, select the check box in the Show Tag column of its row.

Inventory : List of Objects

vRealize Operations Manager discovers objects in your environment for each adapter instance and lists them. From the complete list of all the objects in your environment, you can quickly access and configure any object. For example, you can check if a datastore is connected or providing data, or you can power on a virtual machine.

How the List Works

Objects appear in a data grid. To find a particular object, you can sort a column in the grid or search for a filter word. In addition to sorting and searching, assigning objects to object tags makes it easier to find objects and metrics.

Where You Find the List

In the menu, click **Administration**, then click **Inventory** . The system lists all the objects in your environment.

Inventory List Options

The center pane includes object tag options. The right pane includes toolbar options for all of the objects in your environment.

Table 4-145. 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.
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 4-146. Inventory Toolbar Options

Option	Description
Action	Perform an action on the selected object. Available actions depend on the object type. For example, Power on VM applies to the selected virtual machine. See List of vRealize Operations Manager Actions
Open in external application	If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. For example, Open Virtual Machine in a vSphere Client or Search for VM logs in vRealize Log Insight.
Start Collecting	Turn on data collection for the selected object.
Stop Collecting	Do not collect data for the selected object. When data collection stops, vRealize Operations Manager retains metric data for the object in case data collection starts at a later time.
Perform Multi-Collecting	If an object collects metrics through more than one adapter instance, select the adapter instance or instances for data collection. Does not apply to objects that do not use the adapter instance.
Edit object	Edit the selected object. For example, add or change the maintenance schedule for a virtual machine. If multiple objects of the same type are selected, common identifiers for the object type are editable. For example, change the VM entity name of multiple datastores with a single edit. See Manage Objects Workspace .
Add object	vRealize Operations Manager discovers objects for most adapters. For adapters that do not support autodiscovery for all objects, the objects are manually added. See Manage Objects Workspace .
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.

Table 4-146. Inventory Toolbar Options (continued)

Option	Description
Select All	Select all objects displayed.
Show Detail	Display the Summary tab of the selected object.
Per page	The number of objects to list per page.

Manage Objects Workspace

To collect data from an object, you might need to add an object or edit an existing object in your environment. For example, you might need to add objects for an adapter that does not support autodiscovery, or change the maintenance schedule of an existing object.

Where You Find Manage Objects

In the menu, click **Administration**, then click **Configuration > Inventory** in the left pane. Click the plus sign to add an object or the edit icon to edit the selected object.

Items that appear in the window depend on the object that you are editing. Not all options can be changed.

Table 4-147. 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 4-147. Manage Objects Add or Edit Options (continued)

Options	Description
Collection Interval	<p>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.</p> <p>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.</p>
Dynamic Thresholding	<p>On by default, to enable dynamic thresholding and early warning smart alerts. See vRealize Operations Manager Dynamic Thresholds</p>

Discover Objects Workspace

If vRealize Operations Manager does not discover objects after an adapter instance is configured, use manual discovery. Discovering objects is more efficient than adding objects individually.

Note You use discovery to define objects for embedded adapters. vRealize Operations Manager discovers objects that use external adapters.

Where You Find Discover Objects

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Discover Objects** in the List tool bar.

Discover Objects

The Discoveries section of the `describe.xml` file for the adapter might include parameters for discovery information. The `describe.xml` file is in the `conf` subfolder of the adapter, for example `xyz_adapter3/conf/describe.xml`.

Options	Description
Collector	Collector that vRealize Operations Manager uses to discover objects. Only the vRealize Operations Manager Collector is added during installation.
Adapter Type	Adapter type for the objects to discover.
Adapter Instance	Adapter instance of the selected adapter type.

Options	Description
Discovery Info	Selection depends on the adapter type. For example, for a vCenter adapter, the Discovery Info selection adds an option to discover objects of a particular object type.
Only New Objects	On by default, to omit objects that are already discovered.

Discovery Results List

When you use the Discover Objects feature to manually discover objects in your environment, vRealize Operations Manager lists the objects of the specified object type. You can choose the objects to monitor.

Where You Find Discovery Results

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Discover Objects** in the List tool bar.

After you make selections in the Discover Objects Workspace, click **OK**. With the default setting, vRealize Operations Manager displays only newly discovered objects. See [Discover Objects Workspace](#).

Table 4-148. 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 4-149. 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 4-149. Objects (continued)

Options	Description
Exists	Indicates that the object exists in the vRealize Operations Manager environment.
Collect	When selected, imports the object and starts collecting data. Option is active and selectable for newly discovered objects that do not exist in the vRealize Operations Manager environment.

Manage Maintenance Schedules for Your Object Workspace

You use maintenance mode to take an object offline. Many objects in your environment might be intentionally taken offline. For example, you might deactivate a server to update software. If vRealize Operations Manager collects metrics when the object is offline, it might generate incorrect alerts that affect the data for the object's health. When an object is in maintenance mode, vRealize Operations Manager does not collect metrics from the object and does not generate alerts for it.

How Maintenance Schedules Work

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object into maintenance mode from midnight until 3 a.m. every Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can put an object in maintenance mode or take it out of maintenance mode, even if it has an assigned maintenance schedule.

Where You Find Manage Maintenance Schedules

In the menu, select **Administration**, then click **Configuration > Inventory** in the left pane. Click **Start Maintenance** in the List tool bar.

Table 4-150. Manage Maintenance Schedules Options

Options	Description
I will come back and end maintenance myself.	Maintenance mode starts for the selected object when you click OK . You must manually end maintenance mode for this object.
End maintenance in	Type the number of minutes that the object is in maintenance mode.
End maintenance on	Click the calendar icon, and select the date that maintenance mode ends.

Define Custom Property Workspace

In vRealize Operations Manager, you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign custom properties to any subset of objects irrespective of the adapter kind and resource kind. You can use a mouse click, search filter, or a tag selector to select the correct object.

Where You Find Add/Edit Custom Property

In the menu, select **Administration**, then click **Inventory** in the left pane. Click **Add/Edit Custom Property** in the List tool bar.

Table 4-151. Add/Edit Custom Property

Options	Description
Property Name	Select or enter a property name.
Type	Select the property type from the drop-down menu.
Value	Enter a value for the property.

You can assign the custom properties defined in this page to the Custom Object Groups and New Groups.

For more information, see [Custom Object Groups Workspace to Create a New Group](#).

Inventory : Geographical Map of Objects

vRealize Operations Manager discovers objects in your environment for each adapter. Objects that are assigned a GEO Location tag appear on a geographical map. You can use this map to quickly locate your objects in the world.

How the Geographical Map Works

Objects with the GEO Location tag appear on a map of the world.

- To create a GEO Location tag, see [Manage Object Tags Workspace](#).
- To assign objects to the tag, see [Creating and Assigning Tags](#).

Where You Find the Geographical Map

In the menu, select **Administration**, then navigate to **Configuration > Inventory** in the left pane. Click the **Geographical** tab.

Geographical Map Options

Use the plus sign to zoom in. Use the minus sign to zoom out. Click and drag to pan the map to the left or right.

Managing Custom Object Groups in vRealize Operations Manager

A custom object group is a container that includes one or more objects. vRealize Operations Manager uses custom groups to collect data from the objects in the group, and report on the data collected.

Why Use Custom Object Groups?

You use groups to categorize your objects and have the system collect data from the groups of objects and display the results in dashboards and views according to the way you define the data to appear.

You can create static groups of objects, or dynamic groups with criteria that determine group membership as vRealize Operations Manager discovers and collects data from new objects added to the environment.

vRealize Operations Manager provides commonly used object group types, such as World, Environment, and Licensing. The system uses the object group types to categorize groups of objects. You assign a group type to each group so that you can categorize and organize the groups of objects that you create.

Types of Custom Object Groups

When you create custom groups, you can use rules to apply dynamic membership of objects to the group, or you can manually add the objects to the group. When you add an adapter, the groups associated with the adapter become available in vRealize Operations Manager .

- Dynamic group membership. To dynamically update the membership of objects in a group, define rules when you create a group. vRealize Operations Manager adds objects to the group based on the criteria that you define.
- Mixed membership, which includes dynamic and manual.
- Manual group membership. From the inventory of objects, you select objects to add as members to the group.
- Groups associated with adapters. Each adapter manages the membership of the group. For example, the vCenter Server adapter adds groups such as datastore, host, and network, for the container objects in the vSphere inventory. To modify these groups, you must do so in the adapter.

Administrators of vRealize Operations Manager can set advanced permissions on custom groups. Users who have privileges to create groups can create custom groups of objects and have vRealize Operations Manager apply a policy to each group to collect data from the objects and report the results in dashboards and views.

When you create a custom group, and assign a policy to the group, the system uses the criteria defined in the applied policy to collect data from and analyze the objects in the group. vRealize Operations Manager reports on the status, problems, and recommendations for those objects based on the settings in the policy.

Note Only custom groups defined explicitly by users can be exported from or imported to vRealize Operations Manager. Users are able to export or import multiple custom groups. Once an import function has been executed, the user must check to determine if a policy or policies should be associated with the imported group. Export-import operations are available for user defined (created explicitly by user) custom groups only.

How Policies Help vRealize Operations Manager Report On Object Groups

When you apply a policy to an object group, vRealize Operations Manager uses threshold settings, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that you enabled in the policy to collect data from the objects in the group, and report the results in dashboards and views.

When you create a new object group, you have the option to apply a policy to the group.

- To associate a policy with the custom object group, select the policy in the group creation wizard.
- To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

vRealize Operations Manager applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, vRealize Operations Manager applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, vRealize Operations Manager associates the highest ranking policy with that object.

User Scenario: Creating Custom Object Groups

As a system administrator, you must monitor the capacity for your clusters, hosts, and virtual machines. vRealize Operations Manager monitors them at different service levels to ensure that these objects adhere to the policies established for your IT department, and discovers and monitors new objects added to the environment. You have vRealize Operations Manager apply policies to the object groups to analyze, monitor, and report on the status of their capacity levels.

To have vRealize Operations Manager monitor the capacity levels for your objects to ensure that they adhere to your policies for your service levels, you categorize your objects into Platinum, Gold, and Silver object groups to support the service tiers established.

You create a group type, and create dynamic object groups for each service level. You define membership criteria for each dynamic object group to have vRealize Operations Manager keep the membership of objects current. For each dynamic object group, you assign the group type, and add criteria to maintain membership of your objects in the group. To associate a policy with the custom object group, you can select the policy in the group creation wizard.

Prerequisites

- Know the objects that exist in your environment, and the service levels that they support.
- Understand the policies required to monitor your objects.
- Verify that policies are available to monitor the capacity of your objects.

Procedure

- 1 To create a group type to identify service level monitoring, click **Administration** in the menu, then click **Configuration > Group Types**.

- 2 On the Group Types toolbar, click the plus sign and type **Service Level Capacity** for the group type.

Your group type appears in the list.

- 3 Click **Environment** in the menu, then click the **Custom Groups** tab.

- 4 To create a new object group, click the **plus** sign on the Groups toolbar.

The New Group workspace appears where you define the data and membership criteria for the dynamic group.

- a In the Name text box, type a meaningful name for the object group, such as **Platinum_Objects**.
- b In the **Group Type** drop-down menu, select **Service Level Capacity**.
- c (Optional) In the **Policy** drop-down menu, select your service level policy that has thresholds set to monitor the capacity of your objects.

To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

- d Select the **Keep group membership up to date** check box so that vRealize Operations Manager can discover objects that meet the criteria, and add those objects to the group.
- 5 Define the membership for virtual machines in your new dynamic object group to monitor them as platinum objects.
 - a From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Virtual Machine**.
 - b From the empty drop-down menu for the criteria, select **Metrics**.
 - c From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
 - d From the conditional value drop-down menu, select **is less than**.
 - e From the **Metric value** drop-down menu, type **10**.
 - 6 Define the membership for host systems in your new dynamic object group to monitor them as platinum objects.
 - a Click **Add another criteria set**.
 - b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Host System**.
 - c From the empty drop-down menu for the criteria, select **Metrics**.

- d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
 - e From the conditional value drop-down menu, select **is less than**.
 - f From the **Metric value** drop-down menu, type **100**.
- 7 Define the membership for cluster compute resources in your new dynamic object group.
- a Click **Add another criteria set**.
 - b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Cluster Compute Resources**.
 - c From the empty drop-down menu for the criteria, select **Metrics**.
 - d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **capacityRemaining**.
 - e From the conditional value drop-down menu, select **is less than**.
 - f From the **Metric value** drop-down menu, type **1000**.
 - g Click **Preview** to determine whether objects already match this criteria.
- 8 Click **OK** to save your group.

When you save your new dynamic group, the group appears in the Service Level Capacity folder, and in the list of groups on the **Groups** tab.

- 9 Wait five minutes for vRealize Operations Manager to collect data from the objects in your environment.

Results

vRealize Operations Manager collects data from the cluster compute resources, host systems, and virtual machines in your environment, according to the metrics that you defined in the group and the thresholds defined in the policy that is applied to the group, and displays the results about your objects in dashboards and views.

What to do next

To monitor the capacity levels for your platinum objects, create a dashboard, and add widgets to the dashboard. See [Dashboards](#).

Object Group Types in vRealize Operations Manager

An object group type is an identifier that you apply to a specific group of objects in your environment to categorize them. You can add new group types, and apply them to groups of objects so that vRealize Operations Manager can collect data from the object group and display the results in the dashboards and views.

How the Group Types Work

Use group types to categorize your objects so that the system can apply policies to them to track, and display specific status, such as alerts, workload, faults, risk, and so on.

When you create a new group type, vRealize Operations Manager adds it to the existing list of group types, and creates a new folder with the name of your group type in the Environment Custom Groups list.

When you create a new group of objects, you assign a group type to that group of objects. You add objects from the inventory trees to your custom group, then create your dashboard, add widgets to the dashboard, and configure the widgets to display the data collected from the objects in the group. You can then monitor and manage the objects.

You can apply a group type to a group of objects that you create manually, or to object groups that you cannot modify, such those added by adapters. Each adapter that you add to vRealize Operations Manager adds one or more static groups of objects to group the data received from the adapter sources.

The list of group types appears in the Content area under Group Types. The custom object groups appear in the Environment area under Custom Groups.

Where You Create and Modify a Group Type

To create or modify a group type, click **Administration** in the menu, then **Configuration > Group Types** in the left pane.

Group Type Options

You can add, edit, or delete group types. You cannot edit group types that are created by adapters.

Groups Tab on the Environment Overview Pane

Groups are containers that can contain any number and type of objects in your environment. vRealize Operations Manager collects data from the objects in the group and displays the results in dashboards and views that you define.

How Groups Work

Groups are installed with vRealize Operations Manager, created by an adapter, or created by a user. Based on the group criteria, you can use groups to organize your environment and monitor all objects in the group together. You can also assign policies to groups and make group membership dynamic.

For example, if you have a set of vSphere hosts and you do not want to generate alerts when the host goes into maintenance mode, you can put the vSphere hosts in a group and assign a policy that includes a maintenance schedule setting. During the maintenance period, vRealize Operations Manager ignores any metrics for those objects and does not generate any alerts. After the maintenance period ends, vRealize Operations Manager returns to monitoring the objects and generates alerts if an outage occurs.

Where You Find Custom Groups

To access Custom Groups that you create, click **Environment** on the top menu, then click the **Custom Groups** tab.

Custom Group Options

Click the **ADD** button to add a group. You can only edit, clone, or delete a user-created group. You cannot modify groups installed with vRealize Operations Manager or by an adapter.

You can click the **Horizontal Ellipses** to import or export the custom group. The Groups data grid displays an overview of the state of each group. You can use the All Filters option to sort the custom groups based on Name, ID, Group Type, and Description columns.

To sort the list of custom groups based on columns, click the arrow next to the following columns:

- Name
- Health
- Risk
- Efficiency
- Description
- Members Count

Table 4-152. Group Data Grid Options

Option	Description
Name	Select the group name to display a summary of the group. Select to the right of the name to edit, clone, or delete the group.
Summary	Criticality of the health, risk, and efficiency of any group. Click a group with a red, orange, or yellow criticality to get more details about potential problems with objects in the group.
Members Count	Displays the number of members in the selected group.
Policy	Displays the policy associated with the selected group.
Dynamic Membership	Displays whether the group is static or dynamic. The available options are true and false.
Defined by	Displays who has defined the attributes of the group. The available options are: <ul style="list-style-type: none"> ■ System ■ User Defined ■ Management Pack

Custom Object Groups Workspace

You can create and edit custom groups of objects to have vRealize Operations Manager collect data from the objects and display the results in the dashboards and views so that you can monitor your objects and take action on them when problems occur.

How the Custom Groups Workspace Works

When you create a new object group, you define a meaningful group name, and select the group type. To associate the custom object group with a policy for analysis, you select the policy in the group creation wizard. You can leave the policy selection blank to not associate a policy with the object group. When the policy selection is blank, the custom object group is associated with the policy that is designated as the default policy.

You select the object types, and determine whether membership in the object group is static, dynamic, or a combination of static and dynamic membership.

- To create a static object group, you add objects to the group. You do not include criteria for object membership.
- To create a dynamic object group that vRealize Operations Manager updates based on specific criteria, you select the object type and define membership criteria for the group based on metrics, relationships, and properties.

When you add objects to a custom object group, a new folder appears in the Custom Groups navigation pane on the left, and includes the member objects.

Where You Create and Modify Object Groups

To create or modify static or dynamic object groups, or object groups that have a combination of static and dynamic membership, click **Environment > Custom Groups**. The **Custom Groups** tab displays a list of custom object groups, and the object groups for adapters added to vRealize Operations Manager .

To edit existing groups, select a group and click the edit icon on the **Custom Groups** tab.

Custom Object Groups Workspace to Create a New Group

You can create a new object group, define custom properties, assign a group type and objects to the group. When you create the group, you can assign a policy, or leave the policy selection blank to apply the default policy. vRealize Operations Manager collects data from the objects in the group based on the settings in the policy that is associated with the group. The results appear in the dashboards and views.

Where You Assign Custom Group Type, Policy, and Membership

To assign the group type, policy, and membership, click **Environment**, click **Custom Groups**, and click the plus sign to add a new group. In the New Group workspace, you can define the membership criteria, and select the objects to include or exclude.

To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

Table 4-153. 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.

Table 4-153. New Group Workspace (continued)

Option	Description
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	<p>Defines the criteria for a dynamic object group and has vRealize Operations Manager keep the object membership of the group current.</p> <ul style="list-style-type: none"> ■ 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 4-153. New Group Workspace (continued)

Option	Description
Objects To Always Include pane	<p>Determine which objects to include in the group every time vRealize Operations Manager collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of vRealize Operations Manager , these objects were called a allowlist.</p> <ul style="list-style-type: none"> ■ 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. <ul style="list-style-type: none"> ■ 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. <ul style="list-style-type: none"> ■ Selected objects only. Removes only the selected objects from the list of objects to always include. ■ Selected objects and direct children. Removes the selected objects and the children of the selected objects from the list of objects to always include. ■ Selected objects and all descendants. Removes the selected objects and the descendants of the selected objects from the list of objects to always include.

Table 4-153. New Group Workspace (continued)

Option	Description
Objects To Always Exclude pane	<p>Determine which objects to exclude from the group every time vRealize Operations Manager collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of vRealize Operations Manager , these objects were called a denylist.</p> <ul style="list-style-type: none"> ■ 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. <ul style="list-style-type: none"> ■ Selected objects only. Adds only the selected objects to be permanently excluded from the object group. ■ Selected objects and descendants. Adds the selected objects and the descendants of the selected objects for permanent exclusion from the object group. ■ Objects to always exclude (n) pane. Lists the objects that you add to the exclude list. You must select the check box in the right pane to confirm exclusion of the objects. The number of objects selected for exclusion is reflected by the (n) variable in the title of the pane. ■ Remove button. Removes the objects selected in the right pane from the list of objects to always exclude. <ul style="list-style-type: none"> ■ Selected objects only. Removes only the selected objects from the list of objects to always exclude. ■ Selected objects and direct children. Removes the selected objects and the children of the selected objects from the list of objects to always exclude. ■ Selected objects and all descendants. Removes the selected object and the descendants of the selected objects from the list of objects to always exclude.
Assign Custom Properties	<p>In vRealize Operations Manager , you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign the newly defined custom properties to new groups or existing groups.</p> <ul style="list-style-type: none"> ■ Property Name. Select or specify a name for the custom property. ■ Type. Select the type of custom property from the drop-down menu. <p>The custom property can either be a string or a numeric.</p> <ul style="list-style-type: none"> ■ Inclusion Value. Specify a custom property value, which should be assigned to this custom property when an object is added to the group. ■ Exclusion Value. Specify a custom property value, which should be assigned to this custom property when an object leaves the group. ■ Reset. Resets the custom property to a non-zero value. ■ Remove. Removes the custom property from the group. ■ Add Another Custom Property. Adds another custom property to the group.

Managing Application Groups

An application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. vRealize Operations Manager builds an application to determine how your environment is affected when

one or more components in an application experiences problems, and to monitor the overall health and performance of the application. Object membership in an application is not dynamic. To change the application, you manually modify the objects in the container.

Reasons to Use Applications

vRealize Operations Manager collects data from components in the application and displays the results in a summary dashboard for each application with a real-time analysis for any of the components. If a component experiences problems, you can see where in the application the problems arise, and determine how problems spread to other objects.

Note vRealize Operations Manager provides for calendar periodicity. If your application includes work performed on a specific day of the month, for example, the 15th of the month or the last day of the month, this calendar function identifies the pattern after six cycles of the application. Once the pattern is recognized, the system can forecast accurately into the future. Because the system acquires its information from the input data, you do not have to give any details about how you schedule periodical work.

Applications Tab on the Environment Overview Pane

Applications are groups of related objects in your environment that mimic an application in your business. Use the summary to track the health of objects in the application and help troubleshoot performance issues.

How Applications Work

In vRealize Operations Manager, each application contains one or more tiers and each tier contains one or more objects. The tier is a convenient way to organize objects that perform a specific task in an application. For example, you can group all of your database servers together in a tier.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the application.

Construct an application to view a particular segment of your business. The application shows how the performance of one object affects other objects in the same application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data for your business, you see a yellow, orange, or red status if the application health is degrading. Starting with the application summary dashboard, you can investigate which server is causing or exhibiting the problem.

Where You Find Applications

In the menu, click **Environment**, then click the **Applications** tab.

Applications defined in a previous release of vRealize Operations Manager appear after an upgrade.

Application Options

Select an application to edit or delete, or click the **ADD** button to add an application.

The Applications data grid displays an overview of the state of each application.

Table 4-154. Application Data Grid Options

Option	Description
Name	Select the application name to display a summary of the application. Select to the right of the name to edit or delete the application.
Summary	Criticality of the health, risk, and efficiency of any application. Click an application with a red, orange, or yellow criticality to see more details about potential problems with objects in the application.

User Scenario: Adding an Application

As the system administrator of an online training system, you must monitor components in the Web, application, and database tiers of your environment that can affect the performance of the system. You build an application that groups related objects together in each tier. If a problem occurs with one of the objects, it is reflected in the application display and you can open a summary to investigate the source of the problem further.

In your application, you add the DB-related objects that store data for the training system in a tier, Web-related objects that run the user interface in a tier, and application-related objects that process the data for the training system in a tier. The network tier might not be needed. Use this model to develop your application.

Procedure

- 1 In the menu, click **Environment**, then click **Groups and Applications** in the left pane.
- 2 Click the **Applications** tab and click the **ADD** button.
- 3 Click **Basic n-tier Web App** and click **OK**.
 The Application Management page that appears has two rows. Select objects from the bottom row to populate the tiers in the top row.
- 4 Type a meaningful name such as **Online Training Application** in the Application text box.
- 5 For each of the Web, application and database tiers listed, add the objects to the Tier Objects section.
 - a Select a tier name. This is the tier that you populate.
 - b To the left of the object row, select object tags to filter for objects that have that tag value. Click the tag name once to select the tag from the list and click the tag name again to deselect the tag from the list. If you select multiple tags, objects displayed depend on the values that you select.
 You can also search for the object by name.
 - c To the right of the object row, select the objects to add to the tier.
 - d Drag the objects to the Tier Objects section.
- 6 Click **Save** to save the application.

Results

The new application appears in the list of applications on the Environment Overview Applications page. If any of the components in any of the tiers develops a problem, the application displays a yellow or red status.

What to do next

To investigate the source of the problem, click the application name and see Evaluate Object Information Using Badge Alerts and Summary Tab .

To investigate the source of the problem, click the application name and evaluate the object summary information. See the *vRealize Operations Manager User Guide* .

Add Application

When you add an application to an environment, you select from a list of predefined templates or create your own custom template, to group the objects to monitor in your application.

Where You Find Add Application

In the menu, click **Environment**, then **Groups and Applications > Applications** in the left pane. On the **Applications** tab, click the plus sign.

Add Applications Options

Each predefined template provides you with a list of suggested tiers designed to help you group related objects that perform a specific task in your application. After you select an option, you can alter the selection and number of tiers on the Application Management page.

Option	Description
Basic n-tier Web App	Use this template for any basic application.
Advanced n-tier Web App	Use this template for an application that monitors more physical devices, such as the devices that vRealize Operations Manager discovers when you add a network-related Management Pack or Management Packs.
Legacy non-Web App	Use this template for an application that has no Web-related objects.
Network	Use this template for an application that has only network-related objects.
Custom	Select this option to build your own application topology.

Application Management Dialog Box

You use Application Management to select the objects for your application. The objects you select are grouped in tiers and help you to track the health of your application.

Where You Find Application Management

In the menu, click **Environment**, then click the **Groups and Applications** menu and select **Applications**. On the **Applications** tab, click the plus sign. After you select an application template, click OK.

Application Management Options

At the top of the screen, enter a new application name or use the default name from the Add Application page. The application name must be unique.

Below the name, the page is divided into the tier row and the objects row. On each row, selections in the pane on the left filter the selections in the pane on the right.

The tier row is where you select the tiers to populate with objects to monitor for the application.

Table 4-155. 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 4-156. 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 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.

Note This feature is subject to deprecation review in a future release. Use the editor in the widget itself. Specifically, use the table in the Output Data section.

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 4-157. 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.

Widget Definitions List

A widget is a pane on a dashboard that contains information about configured attributes, resources, applications, or the overall processes in your environment. Widgets can provide a holistic, end-to-end view of the health of all the objects and applications in your enterprise. If your user account has the necessary access rights, you can add and remove widgets from your dashboards.

Table 4-158. Summary of Widgets

Widget Name	Description
Alert List	Shows a list of alerts for the objects that the widget is configured to monitor. If no objects are configured, the list displays all alerts in your environment.
Alert Volume	Shows a trend report for the last seven days of alerts generated for the objects it is configured to monitor.
Anomalies	Shows a chart of the anomalies count for the past 6 hours.
Anomaly Breakdown	Shows the likely root causes for symptoms for a selected resource.
Capacity Remaining	Shows a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.
Container Details	Shows the health and alert counts for each tier in a single selected container.
Container Overview	Shows the overall health and the health of each tier for one or more containers.
Current Policy	Shows the highest priority policy applied to a custom group.

Table 4-158. Summary of Widgets (continued)

Widget Name	Description
Data Collection Results	Shows a list of all supported actions specific for a selected object.
DRS Cluster Settings	Shows the workload of the available clusters and the associated hosts.
Efficiency	Shows the status of the efficiency-related alerts for the objects that it is configured to monitor. Efficiency is based on generated efficiency alerts in your environment.
Environment	Lists the number of resources by object or groups them by object type.
Environment Overview	Shows the performance status of objects in your virtual environment and their relationships. You can click an object to highlight its related objects and double-click an object to view its Resource Detail page.
Environment Status	Shows statistics for the overall monitored environment.
Faults	Shows a list of availability and configuration issues for a selected resource.
Forensics	Shows how often a metric had a particular value, as a percentage of all values, within a given time period. It can also compare percentages for two time periods.
Geo	Shows where your objects are located on a world map, if your configuration assigns values to the Geo Location object tag.
Health	Shows the status of the health-related alerts for the objects that it is configured to monitor. Health is based on generated health alerts in your environment.
Health Chart	Shows health information for selected resources, or all resources that have a selected tag.
Heat Map	Shows a heat map with the performance information for a selected resource.
Mashup Chart	Brings together disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs). This widget is typically used for a container.
Metric Chart	Shows a chart with the workload of the object over time based on the selected metrics.
Metric Picker	Shows a list of available metrics for a selected resource. It works with any widget that can provide resource ID.
Object List	Shows a list of all defined resources.
Object Relationship	Shows the hierarchy tree for the selected object.
Object Relationship (Advanced)	Shows the hierarchy tree for the selected objects. It provides advanced configuration options.
Property List	Shows the properties and their values of an object that you select.
Recommended Actions	Displays recommendations to solve problems in your vCenter Server instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines.
Risk	Shows the status of the risk-related alerts for the objects that it is configured to monitor. Risk is based on generated risk alerts in your environment.
Rolling View Chart	Cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget.

Table 4-158. 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, risk, and efficiency scores for selected resources.
Sparkline Chart	Shows graphs that contain metrics for an object . If all the metrics in the Sparkline Chart widget are for an object that another widget provides, the object name appears at the top right of the widget.
Tag Picker	Lists all defined resource tags.
Text Display	Reads text from a Web page or text file and shows the text in the user interface.
Time Remaining	Shows a chart of the Time Remaining values for a specific resource over the past 7 days.
Top Alerts	Lists the alerts most likely to negatively affect your environment based on the configured alert type and objects.
Top-N	Shows the top or bottom N number metrics or resources in various categories, such as the five applications that have the best or worst health.
Topology Graph	Shows multiple levels of resources between nodes.
View	Shows a defined view depending on the configured resource.
Weather Map	Uses changing colors to show the behavior of a selected metric over time for multiple resources.
Workload	Shows workload information for a selected resource.
Workload Pattern	Shows a historical view of the hourly workload pattern of an object.

For more information about the widgets, see the vRealize Operations Manager help.

Alert List Widget

The Alert List widget is a list of alerts for the objects it is configured to monitor. You can create one or more alert lists in vRealize Operations Manager for objects that you add to your custom dashboards. The widget provides you with a customized list of alerts on objects in your environment.

How the Alert List Widget and Configuration Options Work

You can add the Alert List widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. You edit an Alert List widget after you add it to a dashboard. The changes you make to the options create a custom alert list to meet the needs of the dashboard users.

Criticality	Alert	Triggered On	Created On	Status	Alert Type	Alert Subtype
Yellow	Virtual machine disk I/O write latency	Rima-Demo	2:06 PM	Yellow	Storage	Performance
Orange	Virtual machine disk I/O write latency	11726572_271017...	2:01 PM	Yellow	Storage	Performance
Yellow	Virtual machine disk I/O write latency	VC_60_server1_50	2:01 PM	Yellow	Storage	Performance
Yellow	Virtual machine disk I/O write latency	ESX_6.0_for_VC...	1:56 PM	Yellow	Storage	Performance
Yellow	Virtual machine disk I/O write latency	ESX_5.5_for_VC...	1:56 PM	Yellow	Storage	Performance
Red	Host in a cluster that does not have	evn-lab-esx-38.e...	1:56 PM	Yellow	Virtualiza...	Performance
Yellow	Virtual machine disk I/O write latency	vRealize Operatio...	1:56 PM	Yellow	Storage	Performance
Red	Virtual Machine on a host with BIOS	vRealize Operatio...	1:51 PM	Yellow	Virtualiza...	Performance
Yellow	Virtual machine disk I/O write latency	VA_lib_test_gagi...	1:51 PM	Yellow	Storage	Performance
Yellow	Virtual machine disk I/O write latency	cert-test-client-01	1:51 PM	Yellow	Storage	Performance

Where You Find the Alert List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Alert List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	<p>Actions you can run on the selected alert.</p> <p>For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems.</p>
Reset Interaction	<p>Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.</p> <p>Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards.</p>

Option	Description
Perform Multi-Select Interaction	<p>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.</p> <p>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.</p>
Display Filtering Criteria	Displays the object information on which this widget is based.
Select Date Range	Limits the alerts that appear in the list to the selected date range.
Cancel Alert	<p>Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.</p> <p>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.</p>
Suspend	<p>Suspend an alert for a specified number of minutes.</p> <p>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.</p> <p>The user who suspends the alert becomes the assigned owner.</p>
Take Ownership	<p>As the current user, you make yourself the owner of the alert.</p> <p>You can only take ownership of an alert, you cannot assign ownership.</p>
Release Ownership	Alert is released from all ownership.
Group By	Group alerts by the options in the drop-down menu.
Filter	Locate data in the widget.

Table 4-159. Group By Options

Option	Description
None	Alerts are not sorted into specific groupings.
Time	Group alerts by time triggered. The default.
Criticality	Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the Alert List Widget Data Grid table.
Definition	Group alerts by definition, that is, group like alerts together.
Object Type	Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together.

Alert List Widget Data Grid Options

The data grid provides information on which you can sort and search.

Expand the grouped alerts to view the data grid.

Option	Description
Criticality	<p>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.</p> <p>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.</p>
Alert	Description of the alert.
Triggered On	Name of the object for which the alert was generated.
Created On	Date and time when the alert was generated.
Status	Current state of the alert.
Alert Type	<p>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.</p> <p>The possible values include:</p> <ul style="list-style-type: none"> ■ Application ■ Virtualization/Hypervisor ■ Hardware (OSI) ■ Storage ■ Network

Option	Description
Alert Sub-Type	<p>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.</p> <p>The possible values include:</p> <ul style="list-style-type: none"> ■ Availability ■ Performance ■ Capacity ■ Compliance ■ Configuration
Importance	Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm.

Alert List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Filter	
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.</p> <p>If the objects have an input transformation applied, you select tag values for the transformed objects.</p>


Option	Description
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.
Alert Related	<p>A group of filters limits the alerts that appear in this alert list to those that meet the selected criteria.</p> <p>If the objects on which the alerts are based have an input transformation applied, you define filters for the alerts based on the transformed objects.</p> <p>You can configure the following filters:</p> <ul style="list-style-type: none"> ■ Alert Type. Select the subtype in the type list. This value was assigned when you configured the alert definition. ■ Status. Select one or more alert states to include in the list. ■ Control State. Select one or more control states to include in the list. ■ Criticality. Select one or more levels of criticality. ■ Impact. Select one or more alert badges to include in the list.

Alert Volume Widget

The Alert Volume widget is a trend report for the last seven days of alerts generated for the objects it is configured to monitor in vRealize Operations Manager. You can create one or more alert volume widgets for objects that you add to your dashboards. The alert volume provides you with a customized trend report on objects that helps you identify changes in alert volume, indicating a problem in your environment.

How the Alert Volume Widget and Configuration Options Work

You can add the Alert Volume widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. The changes you make to the options create a custom widget to meet the needs of the dashboard users.



Where You Find the Alert Volume Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Alert Volume Widget Display Options

The Alert Volume widget displays a trend chart, symptoms by criticality, and active alerts.

Option	Description
Trend chart	Volume of critical, immediate, and warning symptoms for the configured objects.
Symptoms by criticality	Number of symptoms for each criticality level.
Active Alerts	Number of active alerts. Alerts can have more than one triggering symptom.

Alert Volume Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.


The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Anomalies Widget

The Anomalies widget displays the anomalies for a resource for the past 6 hours at time intervals you set.

The Anomalies widget shows or hides time periods when the metric violates a threshold that configured. The widget color indicates the criticality of the violation.



Where You Find the Anomalies Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Anomalies Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.


Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Anomaly Breakdown Widget

The Anomaly Breakdown widget shows the likely root causes for symptoms for a selected resource.

How the Anomaly Breakdown Widget and Configuration Options Work



You can add the Anomaly Breakdown widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

Where You Find the Anomaly Breakdown Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Anomaly Breakdown Widget Display Options

The Anomaly Breakdown widget displays scores, volume, and a list of anomaly metrics.

Option	Description
Score	Anomaly value.
Volume	vRealize Operations Manager full set metric count for the selected object in the specified time range.
Anomaly Metrics List	List of alarms for the selected object in the specified time range.

Anomaly Breakdown Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Show Bar Details	If the widget is displaying data for multiple objects, you can select a row and click this button to view the list of alarms for the selected object.
Perform Multiple Interaction	<p>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.</p> <p>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.</p>

Anomaly Breakdown Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

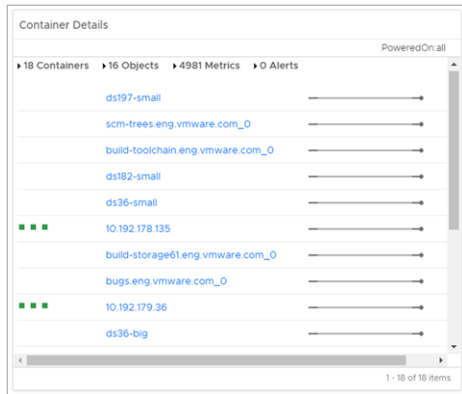
The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Mode	Display a single object or multiple objects.
Show	Select the number of objects to display in multiple objects mode.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.

Container Details Widget

The Container Details widget displays graphs that show a summary of child objects, metrics, and alerts of an object in the inventory.



How the Container Details Widget and Configuration Options Work

The Container Details widget treats objects from the inventory as containers and objects. Containers are objects that contain other objects. The widget lists the containers and shows the number of containers, objects, metrics, and alerts of the observed object. The widget also displays the alerts of each container and an icon links to its child objects. For example, if you select from the inventory a host that contains three objects such as, two virtual machines and one datastore, the Container Details widget displays summary information with three containers, two objects that are the child objects of the two virtual machines, and the number of alerts for the host and the number of metrics for the child objects of the host. The widget also lists each of the three containers, with the number of alerts for each object. Clicking an object in the graph takes you to the object details page. When you point to the icon next to the object, a tool tip shows the name of the related resource and its health. For example, when you point to the icon next to a virtual machine, the tool tip shows a related datastore and its health. Clicking the icon takes you to the object detail page of the related object, which is the datastore following the example.

You edit a container details widget after you add it to a dashboard. You can configure the widget to take information from another widget in the dashboard and to analyze it. When you select **Off** from the Self Provider option and set source and receiver widgets in the **Widget Interactions** menu during editing of the dashboard, the receiver widget shows information about an object that you select from the source widget. For example, you can configure the Container Details widget to display information about an object that you select from the Object Relationship widget in the same dashboard.

Where You Find the Container Details Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Container Details Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Mode	You can change the size of the graph using the Compact or Large buttons.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Capacity Remaining Widget

The Capacity Remaining widget displays a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.

Where You Find the Capacity Remaining Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Capacity Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.


























The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Container Overview Widget

The Container Overview widget gives a graphical presentation of the health, risk, and efficiency of an object or list of objects in the environment.

Container Overview				
Name	Health	Risk	Efficiency	
 v				
 C				
 A				
 v				
 v				
 v				
 1 - 50 of 421 items < 1 2 3 4 5 ... 9 >				

How the Container Overview Widget and Configuration Options Work

The Container Overview widget displays the current status, the status for a previous time period of the health, risk, and the efficiency of an object or list of objects. You can configure the widget to display information for one or more objects that you are interested in when you select the **Object** mode during configuration of the widget. The widget displays information for all objects from an object type or types when you select the **Object Type** mode during configuration of the widget. You can open the object detailed page of each object in the data grid when you click the object.

You edit a container overview widget after you add it to a dashboard. You can configure the widget to display information about an object or to display information about all objects from an object type by using the **Object** or **Object Type** mode. The configuration options change depending on your selection of mode.

Where You Find the Container Overview Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Container Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about other widgets or dashboards.

Option	Description
Perform Multi-Select Interaction	<p>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.</p> <p>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.</p>
Filter	You can filter the objects in the data grid.
Dashboard Navigation	<p>You can explore information from another dashboard.</p> <p>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.</p> <p>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.</p>

Container Overview Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Name	Name of the object
Health	<p>Shows information about the health parameter.</p> <p>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.</p> <p>Last 24 Hours displays the statistic of health parameter for last 24 hours.</p>

Option	Description
Risk	Shows information about the risk parameter. Status displays the badge of the current risk status of an object. You can check the status in a tool tip when you point to the badge. Last Week displays the statistics of the health parameter for the last week.
Efficiency	Shows information about the efficiency parameter. Status displays the badge of the current efficiency status of an object. You can check the status in a tool tip when you point to the badge. Last Week displays statistic of the efficiency parameter for the last week.

Container Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Mode	Use Object to select an object from the environment to observe. Use Object Type to select the type of the objects to observe.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
Object Type	<p>Select an object type in your environment on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add Object Type icon to search for and add an object type. <p>When you search for object types, you can filter the types in the list by selecting a type from the Adapter Type drop-down menu or by using the Filter text box.</p> <ol style="list-style-type: none"> 2 Optionally, select the object type from the list and click the Delete Object Type icon to remove the selected object type.

Current Policy Widget

The Current Policy widget displays the active operational policy that is assigned to your object or object group. vRealize Operations Manager uses the assigned policy to analyze your objects, control the data that is collected from those objects, generate alerts when problems occur, and display the results in the dashboards.

How the Current Policy Widget and Configuration Options Work

You add the Current Policy widget to a dashboard so that you can quickly see which operational policy is applied to an object or object group. To add the widget to a dashboard, you must have access permissions associated with the roles assigned to your user account.

The configuration changes that you make to the widget creates a custom instance of the widget that you use in your dashboard to identify the current policy assigned to an object or object group. When you select an object on the dashboard, the policy applied to the object appears in the Current Policy widget, with an embedded link to the policy details. To display the inherited and local settings for the applied policy, click the link.

Where You Find the Current Policy Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Current Policy Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ 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. <p>For example, to view the policy applied to each object that you select in the Object List widget, select Off for Self Provider.</p>
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>

Data Collection Results Widget

The Data Collection Result widget shows a list of all supported actions specific for a selected object. The widget retrieves data specific to a selected object actions and uses the action framework to run data collection actions.

How the Data Collection Results Widget and Configuration Options Work

You can add the Data Collection Results widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The Data Collection Results widget is a receiver of a resource or metric ID. It can interact with any resource or metric ID that provides widgets such as Object List and Metric Picker. To use the widget, you must have an environment that contains the following items.

- A vCenter Adapter instance
- A vRealize Operations Manager for Horizon View Adapter
- A vRealize Operations Manager for Horizon View Connection Server

You edit a Data Collection Result widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where You Find the Data Collection Results Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Data Collection Results Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Results	Shows all finished and currently running actions for the selected object.
Choose Action	Shows a list with all supported actions specific for the selected object. The selected object is a result of widget interactions.

Data Collection Results Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget updates only when you open the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Config	Specifies self provider choice and selection of a resource instance.
Selected Object	When you select an object, this text box is populated by the object.

Option	Description
Start new data collection on interaction change	Indicates whether to start a new data collection action when the object selection changes in the source widget.
Objects	List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget.
Defaults	Specifies the default data collection action selected for each object type.
Object Types	List of object types in your environment that you can search or sort by column so that you can locate the object type on which you are basing the data that appears in the widget. You can filter the types in the list by selecting a type from the Adapter Type drop-down menu or by using the Filter text box.
Default Data Collection Action	This panel is populated by the object type that you select in the object types list. You can select only one default data collection action for an object type.

DRS Cluster Settings Widget

The DRS Cluster Settings widget displays the workload of the available clusters and the associated hosts. You can change the Distributed Resource Scheduler (DRS) automation rules for each cluster.

How the DRS Cluster Settings Widget and Configuration Options Work

You can view CPU workload and memory workload percentages for each of the clusters. You can view CPU workload and memory workload percentages for each host in the cluster by selecting a cluster in the data grid. The details are displayed in the data grid below. You can set the level of DRS automation and the migration threshold by selecting a cluster and clicking **Cluster Actions > Set DRS Automation**.

DRS Cluster Settings ⌵ ✎ ? 🔍

Name	Datacenter	vCenter	DRS Settings	Migration Threshold	CPU Workload %	Memory Workload %
DRS-Cluster10-001	DRS-Horizon-10	vc_10-27-001-10	✓ Fully Automated	Most Aggressive	<div><div></div></div> ?	<div><div></div></div> ?
DRS-Cluster11-001	DRS-Horizon-11	vc_11-27-001-11	✓ Fully Automated	Default	<div><div></div></div> 21%	<div><div></div></div> 53%
DRS-Cluster12-001	DRS-Horizon-12	vc_12-27-001-12	✓ Fully Automated	Default	<div><div></div></div> 31%	<div><div></div></div> 103%
DRS-Cluster13-001	DRS-Horizon-13	vc_13-27-001-13	✓ Fully Automated	Default	<div><div></div></div> ?	<div><div></div></div> ?
DRS-Cluster14-001	DRS-Horizon-14	vc_14-27-001-14	✓ Fully Automated	Default	<div><div></div></div> ?	<div><div></div></div> ?
DRS-Cluster15-001	DRS-Horizon-15	vc_15-27-001-15	✗ Disabled	--	<div><div></div></div> ?	<div><div></div></div> ?
DRS-Cluster16-001	DRS-Horizon-16	vc_16-27-001-16	✗ Disabled	--	<div><div></div></div> 23%	<div><div></div></div> 51%
DRS-Cluster17-001	DRS-Horizon-17	vc_17-27-001-17	✓ Fully Automated	Default	<div><div></div></div> 13%	<div><div></div></div> 36%
DRS-Cluster18-001	DRS-Horizon-18	vc_18-27-001-18	✗ Disabled	--	<div><div></div></div> 9%	<div><div></div></div> 28%
DRS-Cluster19-001	DRS-Horizon-19	vc_19-27-001-19	✓ Fully Automated	Default	<div><div></div></div> 13%	<div><div></div></div> 93%
DRS-Cluster20-001	DRS-Horizon-20	vc_20-27-001-20	✓ Fully Automated	Default	<div><div></div></div> 16%	<div><div></div></div> 68%
DRS-Cluster21-001	DRS-Horizon-21	vc_21-27-001-21	✓ Fully Automated	Default	<div><div></div></div> 19%	<div><div></div></div> 60%

1 - 13 of 13 items

You edit a DRS Cluster Settings widget after you add it to a dashboard. To configure the widget, click the edit icon at the upper-right corner of the widget window. You can add the DRS Cluster Settings widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The DRS Cluster Settings widget appears on the dashboard named vSphere DRS Cluster Settings, which is provided with vRealize Operations Manager.

Where You Find the DRS Cluster Settings Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

DRS Cluster Settings Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Cluster Actions	Limits the list to actions that match the cluster you select.
Show	The drop-down menu displays the parent vCenter Server instances where the clusters reside. You can also view the data centers under each parent vCenter Server instance. Select a parent vCenter Server to view the workload of the available clusters in the data grid. The default setting displays the clusters across all vCenters.
Filter	Filters the data grid by name, data center, vCenter, DRS settings, and migration threshold.

DRS Cluster Settings Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Name	Displays the names of the clusters in the selected parent vCenter Server instance.
Datacenter	Displays the data centers that belong to each cluster.
vCenter	Displays the parent vCenter Server instance where the cluster resides.
DRS Settings	Displays the level of DRS automation for the cluster. To change the level of DRS automation for the cluster, select Cluster Actions > Set DRS Automation from the toolbar. You can change the automation level by selecting an option from the drop-down menu in the Automation Level column.
Migration Threshold	Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster.
CPU Workload %	Displays the percentage of CPU in GHz available on the cluster.
Memory Workload %	Displays the percentage of memory in GB available on the cluster.

DRS Cluster Settings Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Efficiency Widget

The efficiency widget is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in vRealize Operations Manager usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.

How the Efficiency Widget and Configuration Options Work

You can add the efficiency widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning efficiency alerts generated over time, if the monitored object is a group.
- A trend line displays the efficiency status of the monitored object over time if the object does not provide its resources to any other object, or where no other object depends on the monitored object's resources. For example, if the monitored object is a virtual machine or a distributed switch.
- A pie chart displays the reclaimable, stress, and optimal percentages for the virtual machines that are descendants of the monitored object for all other object types. You use the chart to identify objects in your environment from which you can reclaim resources. For example, if the object is a host or datastore.

If the **Badge Mode** is set to **On**, only the badge appears.

Edit an efficiency widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Efficiency Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Efficiency Widget Display Options

The Efficiency widget displays an efficiency badge. The widget also displays an efficiency trend when not in badge mode.

Option	Description
Efficiency Badge	Status of the objects configured for this instance of the widget. Click the badge to open the Alerts tab for the object that provides data to the widget.
Efficiency Trend	Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the Badge Mode configuration option is off. If the Badge Mode is on, only the badge appears.

Efficiency Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Badge Mode	<p>Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> ■ On. Only the badge appears in the widget. ■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>


Environment Widget

The Environment widget displays the resources for which collects data. You can create one or more lists in vRealize Operations Manager for the resources that you add to your custom dashboards.

How the Environment Widget and Configuration Options Work

The Environment widget lists the number of resources by object or groups them by object type. You can add the Environment widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an Environment widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.



Where You Find the Environment Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.


The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>

Environment Overview Widget

The Environment Overview widget displays the health, risk, and efficiency of resources for a given object from the managed inventory.



How the Environment Overview Widget and Configuration Options Work

You can add the Environment Overview widget to one or more custom dashboards.

The widget displays data for objects from one or several types. The data that the widget displays depends on the object type and category that you selected when you configured the widget.

The objects in the widget are ordered by object type.

The parameters for the health, risk, and efficiency of an object appear in a tool tip when you point to the object.

When you double-click an object on the Environment Overview widget, you can view detailed information for the object.

To use the Environment Overview widget, you must add it to the dashboard and configure the data that appears in the widget. You must select at least one badge and an object. Additionally, you can select an object type.

The Environment Overview widget has basic and advanced configuration options. The basic configuration options are enabled by default.

To use all features of the Environment Overview widget, you must change the default configuration of the widget. Log in to the vRealize Operations Manager machine and set `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. The change is propagated after you use the `service vmware-vcops-web restart` command to restart the UI.

You must use the **Badge** tab to select the badge parameters that the widget shows for each object. You must use the **Config** tab to select an object or object type. To observe a concrete object from the inventory, you can use the **Basic** option. To observe a group of objects or objects from different types, you must use the **Advanced** option.

Where You Find the Environment Overview Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about badges.

Option	Description
Badge	You can select a Health, Risk, or Efficiency badge for objects that appear in the widget. The tool tip of a badge shows the standard name of the badge.
Status	You can filter objects based on their badge status and their state.
Sort	You can sort objects by letter or by number.

Environment Overview Widget Configuration Options


On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Selected Object	Object that is the basis for the widget data. To populate the text box, select Config > Basic and select an object from the list.
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Badge	Defines a parameter to observe. You can select or deselect Health, Risk, and Efficiency parameters using check boxes. Default configuration of the widget selects all badges. Select at least one badge parameter.

Option	Description
Config	<p>Basic</p> <p>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.</p> <hr/> <p>Advanced</p> <p>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.</p> <p>Use the Adapter Type drop-down menu to filter the objects types based on an adapter.</p> <p>You can use the Use vSphere Default button to observe the main vSphere object types.</p> <p>To remove an object type from the list, click Remove Selected next to Use vSphere Default.</p> <p>You can use the Object Type Categories menu to select a group or groups of object types to observe.</p> <p>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.</p> <p>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 <code>skittlesCustomMetricAllowed</code> to <code>true</code> in the <code>web.properties</code> file. The <code>web.properties</code> file is located in the <code>/usr/lib/vmware-vcops/user/conf/web</code> folder.</p>

Environment Status Widget

The Environment Status widget displays the statistics for the overall monitored environment.



How the Environment Status Widget and Configuration Options Work

You customize the output of the widget by choosing a category such as Objects, Metrics, Applications, Alerts, Analytics, and Users. You can filter the data by using the tags tree from **Select which tags to filter** in the configuration window.

You edit an environment status widget after you add it to a dashboard. To configure the widget, click the pencil at the right corner of the widget window. You must select at least one type of information from **OBJECTS, METRICS, APPLICATIONS, ALERTS, ANALYTICS, USERS** categories for the widget to display. By default, the widget displays statistics information about all objects in the inventory. You can use the Select which tags to filter option to filter the information. The widget can interact with other widgets in the dashboard, taking data from them and displaying statistics. For example, you can have a Object List widget, which is the source of the data and an Environment Status widget, which is the destination. If you select objects and perform a multiselection interaction from the Object List widget, the Environment Status widget results are updated based on the selections you made in the Object List.

Where You Find the Environment Status Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Environment Status Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.

Option	Description
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1 , the child objects are the transformed inputs for the widget.
Output Data	
Objects	The widget shows summarized information about the objects in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of resources. For example, if you select Adapter Types > Container from Select which tag to filter and click Objects and Objects Collecting , the widget displays the number of containers and collecting containers.
Metrics	The widget shows summarized information about available metrics. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of metrics.
Applications	The widget shows summarized information about available applications. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of applications.
Alerts	The widget shows summarized information about alerts in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of alerts.
Analytics	The widget shows summarized information about the analytics plug-ins. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of analytics.
Users	The widget shows the number of users defined in vRealize Operations Manager. Select Administration > Access Control > User Accounts .
Output Filter	

Option	Description
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.</p> <p>If the objects have an input transformation applied, you select tag values for the transformed objects.</p>
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Faults Widget

The Faults widget displays detailed information about faults experienced by an object

The Faults widget configuration options are used to customize each instance of the widget that you add to your dashboards.

Where You Find the Faults Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Faults Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Forensics Widget

The Forensics widget shows how often a metric has a particular value as a percentage of all values, within a given time period. It can also compare percentages for two time periods.

How the Forensics Widget and Configuration Options Work

You can add the Forensics widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit the Forensics widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where you Find the Forensics Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Forensics Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Percentile	Indicates how much data is above or below the specific value. For example, it indicates that 90% of the data is more than 4 when a vertical line occurs on the value 4.

Option	Description
Input Data	<p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> 1 Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p>

Geo Widget

If your configuration assigns values to the Geo Location object tag, the geo widget shows where your objects are located on a world map. The geo widget is similar to the **Geographical** tab on the Inventory page.

How the Geo Widget and Configuration Options Work

You can move the map and zoom in or out by using the controls on the map. The icons at each location show the health of each object that has the Geo Location tag value. You can add the geo widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a Geo widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

Where You Find the Geo Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Geo Widget Toolbar Options

Option	Description
Zoom in	Zooms in on the map.
Zoom out	Zooms out on the map.

Geo Widget Configuration Options

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ 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.
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Heatmap Widget

The Heatmap widget contains graphical indicators that display the current value of two selected attributes of objects of tag values that you select. In most cases, you can select only from internally generated attributes that describe the general operation of the objects, such as health or the active anomaly count. When you select a single object, you can select any metric for that object.

How the Heatmap Widget and Configuration Options Work

You can add the Heatmap widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The Heatmap widget has a General mode and an Instance mode. The General mode shows a colored rectangle for each selected resource. In the Instance mode, each rectangle represents a single instance of the selected metric for an object.

You can click a color or the size metric box in the bottom of the Heatmap widget to filter the display of cells in the widget. You can click and drag the color filter to select a range of colors. The Heatmap widget displays cells that match the range of colors.

When you point to a rectangle for an object, the widget shows the resource name, group-by values, the current values of the two tracked attributes, virtual machine details, the metric name, and the value of the color. Click **Show Sparkline** to view the value.

You edit a Heatmap widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Heatmap Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Heatmap Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	Actions you can run on the selected alert. For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems.
Group Zoom	You can roll-up non-significant resources with similar characteristics into groups to obtain only the relevant data among the thousands of resources in the system. The roll-up method improves performance and decreases the memory usage. The roll-up box encompasses the average color and the sum of the sizes of all the resources. You can view all the resources by zooming in the roll-up box.
Show/Hide Text	Show or hide the cell name on the heatmap rectangle.
Show Details	If you configure the Heatmap widget as a provider to another widget, such as the Metric Chart widget, you can double-click a rectangle to select that object for the widget. If the widget is in Metric mode, double-clicking a rectangle selects the resource associated with the metric and provides that resource to the receiving widget. Optionally, you can select a cell from the heatmap and click the Show Details icon to see details about the cell.
Reset Interaction	Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.
Reset Zoom	Resets the heatmap display to fit in the available space.
Heatmap Configuration Drop-down	Select from a list of predefined heatmaps.

Heatmap Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	

Option	Description
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.
Output Data	
Configurations	List of saved heatmap configuration options. You can create a configuration and save it in the list. From the options on the right, you can also delete, clone, and reorder the configurations.
Name	Name of the widget.
Group by	First-level grouping of the objects in the heatmap.
Then by	Second-level grouping of the objects in the heatmap.
Relational Grouping	After you select the Group by and Then by objects, select the Relational Grouping check box to reorganize the grouping of the objects, and to relate the objects selected in the Group by text box with the objects selected in the Then by text box.

Option	Description
Mode	<p>General mode</p> <p>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.</p> <p>Instance mode</p> <p>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.</p>
Object Type	Object that is the basis for the widget data.
Size by	<p>An attribute to set the size of the rectangle for each resource.</p> <p>Resources that have higher values for the Size By attribute have larger areas of the widget display. You can also select fixed-size rectangles. In most cases, the attribute lists include only metrics that vRealize Operations Manager generates. If you select a resource kind, the list shows all the attributes that are defined for the resource kind.</p>
Color by	An attribute to set the color of the rectangle for each resource.
Solid Coloring	Select this option to use solid colors instead of a color gradient. By default, the widget assigns red color for high value, brown color for intermediate value and green color for low value. Click the color box to set a different color for the values. You can add up to seven color thresholds by clicking color range.
Color	<p>Shows the color range for high, intermediate and low values. You can set each color and type minimum and maximum color values in the Min Value and Max Value text boxes. By default, green indicates a low value and red indicates the high end of the value range. You can change the high and low values to any color and set the color to use for the midpoint of the range. You can also set the values to use for either end of the color range, or let vRealize Operations Manager define the colors based on the range of values for the attribute.</p> <p>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.</p>
Output Filter	


Option	Description
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.</p> <p>If the objects have an input transformation applied, you select tag values for the transformed objects.</p>
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Health Widget

The Health widget is the status of the health-related alerts for the objects it is configured to monitor in vRealize Operations Manager. Health alerts usually require immediate attention. You can create one or more health widgets for different objects that you add to your custom dashboards.

How the Health Widget and Configuration Options Work

You can add the Health widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.



The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A trend line displays the health status of the monitored object if the object does not provide its resources to any other object. For example, if the monitored object is a virtual machine or a distributed switch.
- A weather map displays the health of the ancestor and descendant objects of the monitored object for all other object types. For example, if the monitored object is a host that provides CPU and memory to a virtual machine.

If the **Badge Mode** is set to **On**, only the badge appears.

You edit a Health widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

Where You Find the Health Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Health Widget Display Options

The Health widget displays a health badge. The widget also displays a health trend when not in badge mode.

Option	Description
Health Badge	<p>Status of the objects configured for this instance of the widget.</p> <p>Click the badge to open the Alerts tab for the object that provides data to the widget.</p> <p>If the Badge Mode option is off, a health weather map or trend chart appears for the object. Whether the map or chart appears depends on the object type. The health weather map displays tool tips for up to 1000 objects.</p>
Health Trend	<p>Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the Badge Mode configuration option is off. If the Badge Mode is on, only the badge appears.</p>

Heath Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Badge Mode	<p>Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> ■ On. Only the badge appears in the widget. ■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>

Health Chart Widget

The Health Chart widget displays Health, Risk, Efficiency, or custom metric charts for selected objects. You use the widget to compare the status of similar objects based on the same value or name.

How the Health Chart Widget and Configuration Options Work

You can add the Health Chart widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

If the widget is configured to display Health, Risk, or Efficiency, the chart values are based on the generated alerts for the selected alert type for the selected objects.

If the widget is configured to display custom metrics, chart values are based on the metric value for the configured time period.

You edit the Health Chart widget after you add it to the dashboard. The changes you make to the options create a custom widget with the selected charts.

The charts are based either on Health, Risk, or Efficiency alert status, or you can base them on a selected metric. You can include a single object, multiple objects, or all objects of a selected type.

To view the value of the object at a particular time, point your cursor over the chart. A date range and metric value tool tip appear.

A context drop-down menu for each chart can be accessed at the top-right corner after the last metric value.

For each chart, you can view the minimum, maximum, and last metric values. The values are displayed at the top-right corner of each chart. Each of the values is preceded by an appropriate icon of the same color as the state of the metric value.

If there is not enough space to view the metric values, a blue information icon is displayed. Point your cursor over the icon to view the metric value details.

Where You Find the Health Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Health Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Date Controls	<p>Use the date selector to limit the data that appears in each chart to the time period you are examining.</p> <p>Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.</p> <p>Dashboard Time is the default option.</p>

Health Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

Option	Description
Close	Deletes the chart.
Save a snapshot	<p>Creates a PNG file of the current chart. The image is the size that appears on your screen.</p> <p>You can retrieve the file in your browser's download folder.</p>
Save a full screen snapshot	<p>Downloads the current graph image as a full-page PNG file, which you can display or save.</p> <p>You can retrieve the file in your browser's download folder.</p>
Download comma-separated data	<p>Creates a CSV file that includes the data in the current chart.</p> <p>You can retrieve the file in your browser's download folder.</p>
Units	Select the units in which the widget displays data. This option is visible when you select a custom source of data in the widget configuration.

Health Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Order By	<p>Determines how the object charts appear in the widget.</p> <p>You can order them based on value or name, and in ascending or descending order.</p>
Chart Height	Controls the height of all charts. Choose from three possible choices - Small, Medium, Large. Default is Medium.
Pagination number	<p>Number of charts that appears on a page.</p> <p>If you prefer scrolling through the charts, select a higher number. If you prefer to page through the results, select a lower number.</p>
Auto Select First Row	Determines whether to start with the first row of data.

Option	Description
Metric	<p>Determines the source of the data.</p> <ul style="list-style-type: none"> ■ Health, Risk, or Efficiency. The displayed charts are based on one of these alert badges. ■ Custom. The displayed charts are based on the selected metric and use either alert symptom state colors or the selected custom color. You can select a unit for the custom metric from the drop-down menu or choose to allow the widget to automatically pick a unit. <p>If you apply custom colors, enter the value in each box that is the highest or lowest value that should be that color. You can select a unit for the metric.</p>
Metric Unit	Select a unit for the custom metric.
Show	<p>Select one or more of the following items to display in the widget:</p> <ul style="list-style-type: none"> ■ Select Object Name to display the name of the object in the widget. ■ Select Metric Name to display the name of the metric in the widget.
Input Data	
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.
Input Transformation	


Option	Description
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.
Output Filter	
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. If the objects have an input transformation applied, you select tag values for the transformed objects.</p>
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Mashup Chart Widget

The Mashup Chart widget shows disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs).

How the Mashup Chart Widget and Configuration Options Work

The Mashup Chart widget contains charts that show different aspects of the behavior of a selected resource. By default, the charts show data for the past six hours.



The Mashup Chart widget contains the following charts.

- A Health chart for the object, which can include each alert for the specified time period. Click an alert to see more information, or double-click an alert to open the Alert Summary page.
- Metric graphs for any or all the KPIs for any objects listed as a root cause object. For an application, this chart shows the application and any tiers that contain root causes. You can select the KPI to include by selecting **Chart Controls > KPIs** on the widget toolbar. Any shared area on a graph indicates that the KPI violated its threshold during that time period.

The metric graphs reflect up to five levels of resources, including the selected object and four child levels.

You edit a Mashup Chart widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

Where You Find the Mashup Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Mashup Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view.

Option	Description
Filters	Filter data based on criticality, status, and alert type.
Event Filters	Filter based on the type of event such as, change, notification, and fault.
Date Controls	<p>Use the date selector to limit the data that appears in each chart to the time period you are examining.</p> <p>Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.</p> <p>Dashboard Time is the default option.</p>
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate.

Mashup Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.


The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Option	Description
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Metric Chart Widget

You can use the Metric Chart widget to monitor the workload of your objects over time. The widget displays data based on the metrics that you select.



How the Metric Chart Widget and Configuration Options Work

You can add the Metric Chart widget to one or more custom dashboards and configure it to display the workload for your objects. The data that appears in the widget is based on the configured menu items for each widget instance.

You edit the Metric Chart widget after you add it to a dashboard. The changes you make to the menu items create a custom widget with the selected metrics that display the workload on your objects.

To select metrics, you can select an object from the object list, then select the metrics. Or, you can select a tag from the object tag list to limit the object list, then select an object. You can configure multiple charts for the same object or multiple charts for different objects.

To use the metric configuration, which displays a set of metrics that you defined in an XML file, the dashboard and widget configuration must meet the following criteria:

- The dashboard **Widget Interaction** menu items are configured so that another widget provides objects to the target widget. For example, an Object List widget provides the object interaction to a chart widget.
- The widget **Self Provider** options are set to **Off**.
- The custom XML file in the **Metric Configuration** drop-down menu is in the `/usr/lib/vmware-vcops/tools/opscli` directory and has been imported into the global storage using the import command.

Where You Find the Metric Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Metric Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Split Charts	Displays each metric in a separate chart.
Stacked Chart	Consolidates all charts into one chart. This chart is useful for seeing how the total or sum of the metric values vary over time. To view the stacked chart, ensure that the split chart option is turned off.
Dynamic Thresholds	Shows or hides the calculated dynamic threshold values for a 24-hour period.
Show Entire Period Dynamic Thresholds	Shows or hides dynamic thresholds for the entire time period of the graph.
Static Thresholds	Shows or hides the threshold values that have been set for a single metric.
Anomalies	Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below.
Trend Line	Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points.
Show Data Values	Enables the data point tooltips if you switched to a zoom or pan option. Show Data Point Tips must be enabled.

Option	Description
Zoom All Charts	Resizes all the charts that are open in the chart pane based on the area captured when you use the range selector. You can switch between this option and Zoom the View .
Zoom the View	Resizes the current chart when you use the range selector.
Pan	When you are in zoom mode, allows you to drag the enlarged section of the chart so that you can view higher or lower, earlier or later values for the metric.
Zoom to Fit	Resets the chart to fit in the available space.
Remove All	Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts.
Refresh Charts	Reloads the charts with current data.
Date Controls	Opens the date selector. Use the date selector to limit the data that appears in each chart to the time period you are examining. Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. Dashboard Time is the default option.
Generate Dashboard	Saves the current charts as a dashboard.

Metric Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

Option	Description
Close	Deletes the chart.
Save a snapshot	Creates a PNG file of the current chart. The image is the size that appears on your screen. You can retrieve the file in your browser's download folder.
Download comma-separated data	Creates a CSV file that includes the data in the current chart. You can retrieve the file in your browser's download folder.
Save a full screen snapshot	Downloads the current graph image as a full-page PNG file, which you can display or save. You can retrieve the file in your browser's download folder.
Units	You can display the data with dots or as a percentage.
Thresholds	You can choose to show/hide Critical , Immediate , and Warning thresholds in the current chart.

Option	Description
Scales	<p>You can choose a scale for a stacked chart.</p> <ul style="list-style-type: none"> ■ Select Linear to view a chart in which the Y axis scale increases in a linear manner. For example, the Y axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on. ■ Select Logarithmic to view a chart in which the Y axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values. <p>Note If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph.</p> <ul style="list-style-type: none"> ■ Select Combined to view overlapping graphs for the metrics. The chart uses individual scales for each graph instead of using a relative scale, and displays a combined view of the graphs. ■ Select Combined by Unit to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs.
Move Down	Moves the chart down one position.
Move Up	Moves the chart up one position.

You can take the following actions on the Metric Chart graph.

Option	Description
Y Axis	Shows or hides the Y-axis scale.
Chart	Shows or hides the line that connects the data points on the chart.
Data Point Tips	Shows or hides the data point tooltips when you hover the mouse over a data point in the chart.
Zoom by X	Enlarges the selected area on the X axis when you use the range selector in the chart to select a subset of the chart. You can use Zoom by X and Zoom by Y simultaneously.
Zoom by Y	Enlarges the selected area on the Y axis when you use the range selector in the chart to select a subset of the chart. You can use Zoom by X and Zoom by Y simultaneously.
Zoom by Dynamic Thresholds	Resizes the Y axis of the chart so that the highest and the lowest values on the axis are the highest and the lowest values of the dynamic threshold calculated for this metric.
Vertical resize	Resizes the height of a graph in the chart.
Remove icon next to each metric name in a stacked chart	Removes the graph for the metric from the chart.

Metric Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	

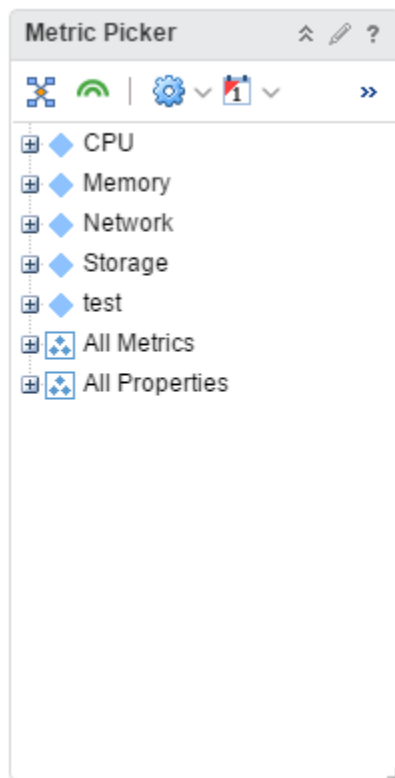
Option	Description
Metrics	<p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select Custom from the Color Method drop-down menu and specify different values for each color, for example 50 for Yellow, 20 for Orange, and 10 for Red.</p> <ol style="list-style-type: none"> Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>

Option	Description
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1 , the child objects are the transformed inputs for the widget.
Output Data	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p> <p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select Custom from the Color Method drop-down menu and specify different values for each color, for example 50 for Yellow, 20 for Orange, and 10 for Red.</p> Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.

Option	Description
Output Filter	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Metric Picker Widget

The Metric Picker widget displays a list of available metrics for a selected object.



How the Metric Picker Widget and Configuration Options Work

With the Metric Picker widget, you can check the list of the object's metrics. To select an object to pick its metrics, you use another widget as a source of data, for example, Topology Graph widget. To set a source widget that is on the same dashboard, you use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a dashboard that contains the source widget. You can also search for objects using tags.

You edit a Metric Picker widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

Where You Find the Metric Picker Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Metric Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Show common metrics	Filter based on common metrics.
Show collecting metrics	Filter based on collecting metrics.
Metrics or Properties	Filter based on metrics or property metrics.
Time Range	Filter based on selected time range.
Search	Search for dashboards, views, and network IP addresses using tags.

Metric Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Action
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Action
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>

Object List Widget

The Object List widget displays a list of the objects available in the environment.

How the Object List Widget and Configuration Options Work

The Object List widget displays a data grid with objects in the inventory. The default configuration of the data grid appears in Object List Widget Options section. You can customize it by adding or removing default columns. You can use the **Additional Column** option to add metrics when you configure the widget.

You edit an Object List widget after you add it to a dashboard. Configuration of the widget enables you to observe parent and child objects. You can configure the widget to display the child objects of an object selected from another widget, for example, another Object List or Object Relationship widget, in the same dashboard.

Click the legend at the bottom of the widget to filter the objects based on threshold. Point your cursor over any of the boxes to view tooltips.

Where You Find the Object List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Action	Selects from a set of actions specific for each object type. To see available actions, select an object from the list of objects and click the toolbar icon to select an action. For example, when you select a datastore object in the graph, you can select Delete Unused Snapshots for Datastore .
Dashboard Navigation	Navigates you to the object. For example, when you select a datastore from the list of objects and click Dashboard Navigation , you can open the datastore in vSphere Web Client.
Reset Grid Sort	Returns the list of resources to its original order.
Reset Interaction	Returns the widget to its initial configured state and undoes any interactions selected in a providing widget. Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards.
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Perform Multi-Select Interaction	If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items. Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to enable the interaction.
Display Filtering Criteria	Displays the object information on which this widget is based.
Page Size	
Filter	Locate data in the widget. You can search for objects or filter the list based on the values of the metrics or properties in the additional columns of the Configuration section.

Object List Widget Data Grid Options

The data grid provides a list of inventory objects on which you can sort and search.

Option	Description
ID	Unique ID for each object in the inventory, randomly generated and produced by vRealize Operations Manager.
Name	Name of the object in the inventory.
Description	Displays the short description of the object given during creation of the object
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 Inventory .
Collection Status	Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the status icon. To manage an adapter instance to start and stop collection of data, in the menu, click Administration , and then in the left pane click Inventory .
Relevance	Displays the user interest on objects based on the number of clicks. The relevance is determined using a system-wide ranking algorithm that rates the object with most clicks as most relevant object.
Internal ID	Unique number that vRealize Operations Manager uses to identify the object internally. For example, the internal ID appears in log files used for troubleshooting.

Object List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Auto Select First Row	Determines whether to start with the first row of data.
Input Data	

Option	Description
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Filter	
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.</p> <p>If the objects have an input transformation applied, you select tag values for the transformed objects.</p>


Option	Description
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.
Additional Columns	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p> <p>Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.</p> <ol style="list-style-type: none"> 1 Click the Add New Metrics icon to add metrics based on object types. The metrics that you add appear in a list in this section. <p>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the Select Object icon to select an object for the object type. Pick metrics of the selected object from the metric tree.</p> <p>For example, you can select the Datacenter object type, click the Select Object icon to display the list of data centers in your environment, and pick metrics of the selected data center.</p> <ol style="list-style-type: none"> 2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click Update.

Object Relationship Widget

The Object Relationship widget displays the hierarchy tree for the selected object. You can create one or more hierarchy trees in vRealize Operations Manager for the selected objects that you add to your custom dashboards.

How the Object Relationship Widget and Configuration Options Work

You can add the Object Relationship widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.



You edit an Object Relationship widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

Where You Find the Object Relationship Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object Relationship Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To be able to navigate to another dashboard, configure the relevant option when you create or edit the dashboard.
Badge	Displays the Health, Risk, or Efficiency alerts on the objects in the relationship map. You can select a badge for objects that appear in the widget. The tool tip of a badge shows the object name, object type, and the name of the selected badge with the value of the badge. You can only select one badge at a time.

Option	Description
Zoom to fit	Resets the chart to fit in the available space.
Pan	Click this icon and click and drag the hierarchy to show different parts of the hierarchy.
Show values on point	Shows or hides the data point tooltips when you hover the mouse over a data point in the chart.
Zoom the view	Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.
Zoom in	Zooms in on the hierarchy.
Zoom out	Zooms out on the hierarchy.
Reset to Initial Object	If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size.
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Show Alerts	Select the resource in the hierarchy and click this icon to show alerts for the resource. Alerts appear in a pop-up window. You can double-click an alert to view its Alert Summary page.

Object Relationship Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Auto Zoom to Fixed Node Size	<p>You can configure a fixed zoom level for object icons in the widget display.</p> <p>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.</p>
Node Size	<p>You can set the fixed zoom level at which the object icons display. Enter the size of the icon in pixels.</p> <p>The widget shows object icons at the pixel size that you configure.</p>
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Object Relationship (Advanced) Widget

The Object Relationship (Advanced) widget displays a graph or tree view that depicts the parent-child relationship of the selected object. It provides advanced configuration options. You can create a graph or tree view in vRealize Operations Manager for the selected objects that you add to your custom dashboards.

How the Object Relationship (Advanced) Widget and Configuration Options Work

You can add the **Object Relationship (Advanced)** widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an **Object Relationship (Advanced)** widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

You can double-click any object in the graph or tree view and see the specific parent-child objects for the focus object. When you double-click the object again, you see the original graph or tree view. If you point your cursor over an object icon, you see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object.

Where You Find the Object Relationship (Advanced) Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Object Relationship (Advanced) Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Options	Description
Dashboard Navigation	You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To navigate to another dashboard, configure the relevant option when you create or edit the dashboard.
Reset to Initial Object	If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.
View Tree/View graph	Displays a tree or graph view of the relationships.
Vertical/Horizontal	Displays a vertical or horizontal view of the graph or tree view.
Hide Text/Show Text	Hides or displays the object names.
Standard View/Fit View	The Standard View option fixes the view to a specific zoom level The Fit View option adjusts the graph or tree view to fit the screen.
Group Items/Ungroup Items	Groups by objects types. You can view further details by double-clicking on the object. You can also choose to display the graph or tree view without grouping the object types.
Path Exploration	Displays the relative relationship path between two selected objects on the graph or tree view. To highlight the path, click the Path Exploration icon and then select the two objects from the graph or tree view.
Layers	<ul style="list-style-type: none"> ■ Parent/Child: Displays a graph or tree view of the parent and child relationship for the specific object selected. ■ Custom: Indicates the relationship between the objects that are part of the custom relationship. These objects have a connection via the selected custom relationship.
Quick Filter	Enter the name of an object that you want to see in the graph or tree view.

Object Relationship (Advanced) Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Name	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Parents Depth	Select the depth of parent objects to be displayed.
Children Depth	Select the depth of child objects to be displayed.
Inventory trees	Select an existing predefined traversal spec for the initial object relationship graph or tree view.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.
Output Filter	

Option	Description
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Property List Widget

You can use the Property List widget to view the properties of objects and their values.

How the Property List Widget and Configuration Options Work

To observe the properties of objects in the Property List widget, you can select object property metrics when you configure the widget itself (Self Provider mode enabled). Alternatively, you can select objects or object property metrics from another widget (Self Provider mode disabled). You can also view a default or custom set of properties by selecting a preconfigured XML file in the Metric Configuration drop-down menu of the widget configuration window.

You edit a Property List widget after you add it to a dashboard. You can configure a widget to receive data from another widget by selecting **Off** for Self Provider mode. When the widget is not in Self Provider mode, it displays a set of predefined properties and their values of an object that you select on the source widget. For example, you can select a host on a Topology widget and observe its properties in the Property List widget. To configure the Property List as a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To configure a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

Where You Find the Property List Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Property List Widget Data Grid Options

The data grid provides information on which you can sort and search.

Option	Description
Object Name	Name of the object, whose properties you observe. You can sort the properties by object name. To open the Object Details page, click an object name.
Property Name	Name of the property. You can sort the properties by property name.
Value	Value of the property. You can sort the properties by value.

Property List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Visual Theme	<p>Select a predefined visual style for each instance of the widget. The options are: Original and Compact.</p>
Show Metric Full Name	<p>You can choose to view the full name of the metrics. The options are: On and Off.</p>
Input Data	
Metrics	<p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the Unit drop-down menu, and click Update.</p>

Option	Description
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Data	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p> <ol style="list-style-type: none"> 1 Click the Add New Metrics icon to add metrics based on object types. The metrics that you add appear in a list in this section. <p>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the Select Object icon to select an object for the object type. Pick metrics of the selected object from the metric tree.</p> <p>For example, you can select the Datacenter object type, click the Select Object icon to display the list of data centers in your environment, and pick metrics of the selected data center.</p> <ol style="list-style-type: none"> 2 Optionally, you can define measurement units for the metrics and properties in the list. Double-click a metric or properties box in the list, select a measurement unit in the Unit drop-down menu, and click Update. 3 You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.

Option	Description
Output Filter	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Recommended Actions Widget

The Recommended Actions widget displays recommendations to solve problems in your vCenter Server instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines.

How the Recommended Actions Widget and Configuration Options Work

The Recommended Actions widget appears on the Home dashboard, and displays the health status for the objects in your vCenter Server instance. At a glance, you can see how many objects are in a critical state, and how many objects need immediate attention.

From the Recommended Actions widget, you can focus in on problems further by, for example, clicking an object where the alerts triggered, and by clicking an individual alert.

You can edit the Recommended Actions widget on the Home dashboard, or on another dashboard where you add the widget. With the widget configuration options, you can assign a new name to the widget, set the refresh content, and set the refresh interval.

The Recommended Actions widget includes a selection bar, a summary pane, a toolbar for the data grid, and alert information for your objects in a data grid.

Where You Find the Recommended Actions Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Recommended Actions Widget Selection Bar and Summary Pane

Option	Description
Scope	Allows you to select an instance of vCenter Server, and a data center in that instance.
Object tabs	Displays the object types with the number of objects affected in parentheses. You can display the actions for virtual machines, host systems, clusters, vCenter Server instances, and datastores.
Badge	<p>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.</p> <ul style="list-style-type: none"> ■ Health Status. With the Health badge selected, displays the number of affected objects and a summary of their health based on the alerts that triggered on the object. Lists the objects that have the worst health, and the number of alerts that triggered on each object. ■ Risk Status. With the Risk badge selected, displays the number of affected objects and a summary of their risk based on the alerts that triggered on the object. Lists the objects that have the highest, and the number of alerts that triggered on each object. ■ Efficiency Status. With the Efficiency badge selected, displays the number of affected objects. Lists the objects that have the lowest efficiency based on the alerts that triggered on the object, and the number of alerts that triggered on each object.
Search filter	Narrows the scope of the objects that appear. Enter a character or a number to search and display an object. When a filter is active, the name of the filter appears below the Search filter text box.

Recommended Actions Widget Toolbar Options

The toolbar allows you to address an alert, and to filter the alert list.

Option	Description
Cancel Alert	<p>Cancels the selected alert.</p> <p>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.</p>
Suspend	<p>Suspends an alert for a specified number of minutes.</p> <p>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.</p> <p>The user who suspends the alert becomes the assigned owner.</p>
All Filters	Narrows the search to one of the available filter types. For example, you can display all alerts that are related to the Compliance Alert Subtype.

Recommended Actions Widget Data Grid Options

The data grid displays the alerts that triggered on your objects. To resolve the problems indicated by the alerts, you can link to the alerts and the objects on which the alerts triggered.

For more information, see [Triggered Alerts](#).

Option	Description
Criticality	<p>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.</p> <p>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.</p>
Actionable	When an alert has an associated action, you can run the action on the object to resolve the alert.
Suggested Fix	<p>Describes the recommendation to resolve the problem. For example, for Compliance alerts, the recommendation instructs you to use the <i>vSphere Hardening Guide</i> to resolve the problem. You can find the <i>vSphere Hardening Guides</i> at http://www.vmware.com/security/hardening-guides.html.</p> <p>You can view other available recommendations and their associated actions, if any, to resolve the problem when you click the drop-down menu.</p>
Name	<p>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.</p> <p>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object.</p>
Alert	<p>Name of the alert definition that generated the alert.</p> <p>Click the alert name to view the alert details tabs where you can begin troubleshooting the alert.</p>
Alert Type	Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/Hypervisor, Hardware, Storage, and Network.
Alert Subtype	Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration.
Time	Date and time that the alert triggered.
Alert ID	Unique identification for the alert. This column is hidden by default.

Recommended Actions Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Risk Widget

The risk widget is the status of the risk-related alerts for the objects it is configured to monitor. Risk alerts in vRealize Operations Manager usually indicate that you should investigate problems in the near future. You can create one or more risk widgets for objects that you add to your custom dashboards.

How the Risk Widget and Configuration Options Work

You can add the risk widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.


The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the Badge Mode configuration option is set to Off, the badge and a chart appear. The type of chart depends on the object type that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning risk alerts generated over time, if the monitored object is a group.
- A trend line displays the risk status of the monitored object for all other object types.

If the Badge Mode is set to On, only the badge appears.

You edit a risk widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.



Where You Find the Risk Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Risk Widget Display Options

The Risk Widget displays a risk badge. The widget also displays a risk trend chart when not in badge mode.

Option	Description
Risk Badge	Status of the objects configured for this instance of the widget. Click the badge to open the Alerts tab for the object that provides data to the widget.
Risk Trend	Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the Badge Mode configuration option is off. If the Badge Mode is on, only the badge appears.

Risk Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.


The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Badge Mode	<p>Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> ■ On. Only the badge appears in the widget. ■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Rolling View Chart Widget

The Rolling View Chart widget cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget.



How the Rolling View Chart Widget and Configuration Options Work

The Rolling View Chart widget shows a full chart for one selected metric at a time. Miniature graphs for the other selected metrics appear at the bottom of the widget. You can click a miniature graph to see the full graph for that metric, or set the widget to rotate through all selected metrics at an interval that you define. The key in the graph indicates the maximum and minimum points on the line chart.

You edit a Rolling View Chart widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

Where You Find the Rolling View Chart Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Rolling View Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Trend Line	Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points.
Dynamic Thresholds	Shows or hides the calculated dynamic threshold values for a 24-hour period.
Show Entire Period Dynamic Thresholds	Shows or hides dynamic thresholds for the entire time period of the graph.

Option	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. Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. Dashboard Time is the default option.

Rolling View Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none">■ On. You define the objects for which data appears in the widget.■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Auto Transition Interval	<p>Time interval for a switch between charts in the widget.</p>
Input Data	

Option	Description
Metrics	<p data-bbox="810 237 1406 294">Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"><li data-bbox="810 306 1406 426">1 Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. The metric tree shows common metrics for several objects when you click the Show common metrics icon. While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<li data-bbox="810 884 1406 968">2 Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. Click the Select All icon to select all the metrics in the list. Click the Clear Selection icon to clear your selection of metrics in the list. <p data-bbox="810 1142 1406 1262">You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the Unit drop-down menu, and click Update.</p>

Option	Description
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Data	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p>

Option	Description
	<p>Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics based on object types. The metrics that you add appear in a list in this section. <p>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the Select Object icon to select an object for the object type. Pick metrics of the selected object from the metric tree.</p> <p>For example, you can select the Datacenter object type, click the Select Object icon to display the list of data centers in your environment, and pick metrics of the selected data center.</p> <ol style="list-style-type: none"> Optionally, you can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the Unit drop-down menu, and click Update.
Output Filter	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> In the first drop-down menu, select an object type. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. In the drop-down menus and text boxes that appear, select or enter values to filter the objects. To add more filter criteria, click Add. To add another filter criteria set, click Add another criteria set.

Scoreboard Widget

The Scoreboard widget shows the current value for each metric of objects that you select.

How the Scoreboard Widget and Configuration Options Work

Each metric appears in a separate box. The value of the metric determines the color of the box. You define the ranges for each color when you edit the widget. You can customize the widget to use a sparkline chart to show the trend of changes of each metric. If you point to a box, the widget shows the source object and metric data. Icons in the box indicate the level of criticality.

You edit a Scoreboard widget after you add it to a dashboard. The widget can display metrics of the objects selected during editing of the widget or selected on another widget. When the Scoreboard widget is not in Self Provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. It shows 10 predefined metrics if you do not select an XML file or if the type of the selected object is not defined in the XML file.

For example, you can configure the Scoreboard widget to use the sample Scoreboard metric configuration and to receive objects from the Topology Graph widget. When you select a host on a Topology Graph widget, the Scoreboard widget shows the workload, memory, and CPU usage of the host.

To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

Where You Find the Scoreboard Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Scoreboard Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p>
Round Decimals	Select the number of decimal places to round the scores that the widget displays.
Box Columns	Select the number of columns that appear in the widget.
Layout Mode	Select a Fixed Size or Fixed View layout.
Fixed Size Fixed View	Use these options to customize the size of the box for each object.
Old metric values	<p>Select Show if you want the widget to show the previous value of the metric, if the current value is not available.</p> <p>Select Hide to hide the previous value of the metric, if the current value is not available.</p>
Visual Theme	Select a predefined visual style for each instance of the widget.
Max Scores Count	Use these menus to customize the format of the scores that the widget displays.

Option	Description
Show	<p>Select one or more of the following items to display in the widget:</p> <ul style="list-style-type: none">■ Select Object Name to display the name of the object in the widget.■ Select Metric Name to display the name of the metric in the widget.■ Select Metric Unit to display the metric unit in the widget.■ Select Sparkline to display the Sparkline chart for each metric.
Period Length	Select a length of time for the statistic information that the sparkline chart displays.
Show DT	Select an option to show or hide the dynamic threshold for the sparkline chart.
Input Data	

Option	Description
Metrics	<p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when</p>

Option	Description
	<p>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.</p> <p>You can use the Link to option to add links to external and internal pages. Internal links open in the same tab. External links open in a new tab. Examples of external links are URLs whose hostname does not match with the current vRealize Operations Manager instance hostname. Internal links are URLs whose hostname matches the current vRealize Operations Manager instance hostname or starts with <i>index.action</i>.</p> <ol style="list-style-type: none"> 2 Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Data	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p>

Option	Description
	<p>Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics based on object types. The metrics that you add appear in a list in this section. <p>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the Select Object icon to select an object for the object type. Pick metrics of the selected object from the metric tree.</p> <p>For example, you can select the Datacenter object type, click the Select Object icon to display the list of data centers in your environment, and pick metrics of the selected data center.</p> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select Custom from the Color Method drop-down menu and specify different values for each color, for example 50 for Yellow, 20 for Orange, and 10 for Red.</p>

Option	Description
	<p>You can use the Link to option to add links to external and internal pages. Internal links open in the same tab. External links will open in a new tab. Examples of external links are URLs whose hostname does not match with the current vRealize Operations Manager instance hostname. Internal links are URLs whose hostname matches the current vRealize Operations Manager instance hostname or starts with <i>index.action</i>.</p> <ol style="list-style-type: none"> 2 Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.
Output Filter	
	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.


Scoreboard Health Widget

The Scoreboard Health widget displays color-coded health, risk, efficiency, and custom metrics scores for objects that you select.

How the Scoreboard Health Widget and Configuration Options Work

The icons for each object are color coded to give a quick indication of the state of the object. You can configure the widget to display the scores of common or specific metrics of the object. You can use the symptom state color code or you can define your criteria to color the images. If you configure the widget to show the metric for objects that do not have this metric, those objects have blue icons.

You can double-click an object icon to show the Object Detail page for the object. When you point to the icon, a tool tip shows the name of the object and the name of the metric.



You edit a Scoreboard Health widget after you add it to a dashboard. To configure the widget, click the pencil at the upper-right corner of the widget window. The widget can display metrics of the objects that you select when you edit the widget, or that you select on another widget. For example, you can configure the widget to show the CPU workload of an object that you select on the Topology Graph widget. To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

Where You Find the Scoreboard Health Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Scoreboard Health Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.


The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Image Type	Select an image type for the metrics.
Metric	Select the default or custom metric.
Pick Metric	<p>Active only when you select Custom from the Metric menu.</p> <p>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.</p> <p>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.</p>
Use Symptom state to color chart	Select to use the default criteria to color the image.
Custom ranges	Use to define custom criteria to color the image. You can define a range for each color.

Option	Description
Input Data	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>

Sparkline Chart Widget

The Sparkline Chart widget displays graphs that contain metrics for an object in vRealize Operations Manager. You can use vRealize Operations Manager to create one or more graphs that contain metrics for objects that you add to your custom dashboards.



How the Sparkline Chart Widget and Configurations Options Work

If the metrics in the Sparkline Chart are for an object that another widget provides, the object name appears at the top right of the widget. If you select a metric when you edit the widget configuration, the widget uses the metric and its corresponding object as the source for dashboard interactions. The line in the graphs represents the average value of the selected metric for the specified time period. The boxed area in the graph represents the dynamic threshold of the metric.

Point to a graph in the Sparkline Chart widget to view the value of a metric in the form of a tool tip. You can also view the maximum and minimum values on a graph. The values are displayed as orange dots.

You can add the Sparkline Chart widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

Where You Find the Sparkline Chart Widget

The widget might be included on any of your custom dashboards. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

Sparkline Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Option	Description
Dashboard Navigation	You can navigate to another dashboard when the object you select is also available in the dashboard to which you want to navigate.
Refresh	Refreshes the widget data.
Time Range	Select the range for the time period to show on the graphs. You can select a period from the default time range list or select start and end dates and times. Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. Dashboard Time is the default option.
Remove All	Removes all graphs.

Sparkline Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	Enable or disable the automatic refreshing of the data in this widget. If not enabled, the widget is updated only when the dashboard is opened or when you click the Refresh button on the widget in the dashboard.
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ 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	<p>You can view the name of the object before the metric name in the Sparkline Chart widget.</p> <ul style="list-style-type: none"> ■ On. Displays the name of the object before the metric name in the widget. ■ Off. Does not display the name of the object in the widget.
Column Sequence	<p>Select the order in which to display the information.</p> <ul style="list-style-type: none"> ■ Graph First. The metric graph appears in the first column in the widget display. ■ Label First. The metric label appears in the first column in the widget display.
Show DT	<p>Select an option to show or hide the dynamic threshold for the sparkline chart.</p>
Input Data	

Option	Description
Metrics	<p>Select metrics on which you want to base the widget data. You can select an object and pick its metrics.</p> <ol style="list-style-type: none"> Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section. <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> Optionally, select metrics from the list and click the Remove Selected Metrics icon to remove the selected metrics. <p>Click the Select All icon to select all the metrics in the list.</p> <p>Click the Clear Selection icon to clear your selection of metrics in the list.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when you</p>

Option	Description
	<p>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.</p> <ol style="list-style-type: none"> 2 Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Data	
Empty drop-down menu	<p>Specifies a list with attributes to display.</p>

Option	Description
	<p>Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.</p> <p>Click the Add New Metrics icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.</p> <p>The metric tree shows common metrics for several objects when you click the Show common metrics icon.</p> <p>While selecting objects for which you want to pick metrics, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <p>Optionally, you can customize a metric and apply the customization to other metrics in the list.</p> <ol style="list-style-type: none"> Double-click a metric box in the list to customize the metric and click Update. <p>You can use the Box Label text box to customize the label of a metric box.</p> <p>You can use the Unit text box to define a measurement unit of each metric.</p> <p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. You can also set coloring by symptom definition. If you do not want to use color, select None.</p> <p>For example, to view the remaining memory capacity of a VM, select Virtual Machine as an object type, expand the Memory from the metric tree and double-click Capacity Remaining(%). Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select Custom from the Color Method drop-down menu and specify different values for each color, for example 50 for Yellow, 20 for Orange, and 10 for Red.</p> Select a metric and click the Apply to All icon to apply the customization for the selected metric to all the metrics in the list.

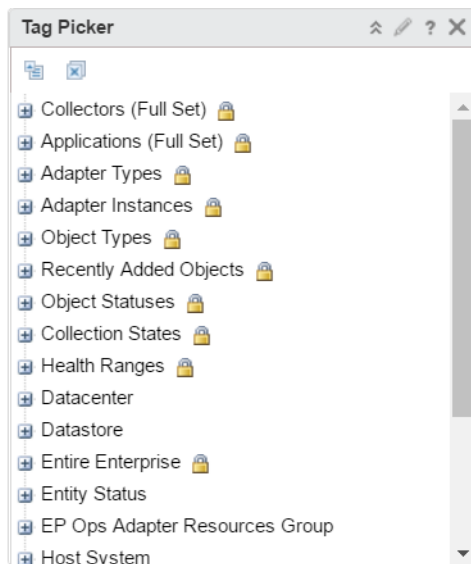
Option	Description
Output Filter	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Tag Picker Widget

The Tag Picker widget lists all available object tags.

How the Tag Picker Widget and Configuration Options Work

With the Tag Picker widget, you can check the list of the object tags. You can use the widget to filter the information that another widget shows. You can select one or more tags from the object tree or search for tags, and the destination widget displays information about the objects with this tag. For example, you can select **Object Types > Virtual Machine** on the Tag Picker widget to observe statistic information about the VMs on the Environment Status widget.



You edit a Tag Picker widget after you add it to a dashboard. To configure the widget, click the pencil in the upper right of the widget window. You can configure the Tag Picker widget to send information to another widget on the same dashboard or on another dashboard. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard. You can configure two Tag Picker widgets to interact when they are on different dashboards.

Where You Find the Tag Picker Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Tag Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Collapse All	Close all expanded tags and tag values.
Deselect All	Remove all filtering and view all objects in the widget.
Tag Picker	Select an object from your environment.
Dashboard Navigation	<p>Note Appears on the source widget and when the destination widget is on another dashboard.</p> <p>Use to explore the information on another dashboard.</p>

Tag Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.

Text Display Widget

You can use the Text Display widget to show text in the user interface. The text appears in the Text Display widget on the dashboard.

The Text Display widget can read text from a Web page or text file. You specify the URL of the Web page or the name of the text file when you configure the Text widget. To use the Text Display widget to read text files you must set a property in the *web.properties* file to specify the root folder that contains the file.

You can enter content in the Text Display widget in plain text or rich text format based on the view mode that you configure. Configure the Text Display widget in HTML view mode to display content in rich text format. Configure the Text Display widget in Text mode to display content in plain text format.

The Text Display widget can display websites that use the HTTPS protocol. The behavior of the Text Display widget with websites that use HTTP, depends on the individual settings of the websites.

Note If the webpage that you are linking to has **X-Frame-Options** set to **sameorigin**, which denies rendering a page in an iframe, the Text Display widget cannot display the contents of the webpage.

How the Text Display Widget Configuration Options Work

You can configure the widget in the Text view mode or HTML view mode. In the HTML view mode, you can click **Edit** in the widget and use the rich text editor to add content.

If you configure the widget to use Text view mode, you can specify the path to the directory that contains the files to read or you can provide a URL. The content in the URL will be shown as text. If you do not specify a URL or text file, you can add content in the widget. Double-click the widget and enter content in plain text.

You can also use command-line interface (CLI) commands to add file content to the Text Display widget.

- To view a list of parameters, run the `file -h|import|export|delete|list txtwidget` command.
- To import text or HTML content, run the `import txtwidget input-file [--title title] [--force]` command.
- To export the content to the file, run the `export txtwidget all|title[{,title}] [output-dir]` command.
- To delete imported content, run the `delete txtwidget all|title[{,title}]` command.
- To view the titles of the content, run the `list txtwidget` command.

Where You Find the Text Display Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Text Display Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.


Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
View mode	Display text in text or rich text format. You can configure the widget in HTML view mode only when the URL and File fields are blank.
URL	Enter the URL.
File	<p>Navigate to the file that contains the source text file by clicking the Browse button.</p> <p>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.</p>
Test	Validates the correctness of the text file or URL that you enter.

Time Remaining Widget

The Time Remaining widget displays how much time remains before the resources of the object are exhausted.

vRealize Operations Manager calculates the percentage by object type based on historical data for the pattern of use for the object type. You can use the time remaining percentage to plan provisioning of physical or virtual resources for the object or rebalance the workload in your virtual infrastructure.



Where You Find the Time Remaining Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Time Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	<p>Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.</p>


Top Alerts Widget

Top alerts are the alerts with the greatest significance on the objects it is configured to monitor in vRealize Operations Manager. These are the alerts most likely to negatively affect your environment and you should evaluate and address them.

How the Top Alerts Widget and Configuration Options Work

You can add the top alerts widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a top alerts widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.



Where You Find the Top Alerts Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Top Alerts Widget Display Options

The Top Alerts widget includes the short description of alerts configured for the widget. The alert name opens a secondary window from which you can link to the alert details. In the alert details, you can begin resolving the alerts.

Option	Description
Alert name	Name of the generated alert. Click the name to open the alert details.
Alert description	Number of affected objects, and the number of recommendations and the best recommendation to resolve the alert.

Top Alerts Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Impact Badge	<p>Select the badge for which you want alerts to appear.</p> <p>The affected badge is configured when you configure the alert definition.</p>
Number of Alerts	Select the maximum number of alerts to display in the widget.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.
Input Transformation	
Relationship	Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1 , the child objects are the transformed inputs for the widget.

Top-N Widget

The Top-N widget displays the top n results from analysis of an object or objects that you select.

How the Top-N Widget and Configuration Options Work

You can select an object when you configure the Top-N widget or you can select an object on another widget. The widget shows an analysis of the applications, alerts, and metrics of an object and its child objects depending on how you configure the widget. The widget can show an analysis of the current values or values over a period of time. You can receive detailed information about each object on the widget. When you double-click an object, the Object Detail page appears.

You can configure a widget to receive data from another widget by selecting **Off** for Self Provider. You can configure a widget to display results from analysis of an object that you select on the source widget.

For example, you can select a host on a Topology widget and observe the metric analysis of the virtual machines on the host. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

Where You Find the Top-N Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Top-N Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

Icon	Description
Dashboard Navigation	Takes you to a predefined object. For example, when you select a datastore from the data grid and click Dashboard Navigation , you can open the datastore in the vSphere Web Client.
Select Date Range	Limits the alerts that appear in the list to the selected date range. Select Dashboard Time to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.
Object details	Select an object and click this icon to show the Object Detail page for the object.
Display Filtering Criteria	Shows the filtering settings for the widget in a pop-up window.

Top-N Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Redraw Rate	Set the redraw rate.
Bars Count	Select the number of top results.
Round Decimals	Select the number of decimals to round the scores displayed in the widget.
Filter old metrics	Select or deselect whether the analysis includes old metric values.

Option	Description
Application Health and Performance	<ul style="list-style-type: none"> ■ 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. <p>Select the criteria for analysis of the objects.</p>
Alert Analysis	Select the criteria for analysis of the alerts.
Metric Analysis	<p>If you select this option, you must select a metric in the Output Data section.</p> <ul style="list-style-type: none"> ■ 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. <p>Select the criteria for analysis of the metric that you select from the metric tree.</p>
Input Data	


Option	Description
Objects	<p>Select objects on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add New Objects icon and select objects in the pop-up window. The selected objects appear in a list in this section. <p>While selecting objects, you can use the Filter text box to search for objects. You can also expand the Tag Filter pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.</p> <ol style="list-style-type: none"> 2 Optionally, select objects from the list and click the Remove Selected Objects icon to remove the selected objects. <p>Click the Select All icon to select all the objects in the list.</p> <p>Click the Clear Selection icon to clear your selection of objects in the list.</p>
All	<p>If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data.</p>
Input Transformation	
Relationship	<p>Transform the input for the widget based on the relationship of the objects. For example, if you select the Children check box and a Depth of 1, the child objects are the transformed inputs for the widget.</p>
Output Data	
	<p>Select an object type in your environment on which you want to base the widget data.</p> <ol style="list-style-type: none"> 1 Click the Add Object Type icon to search for and add an object type. <p>When you search for object types, you can filter the types in the list by selecting a type from the Adapter Type drop-down menu or by using the Filter text box.</p> <ol style="list-style-type: none"> 2 Optionally, select the object type from the list and click the Delete Object Type icon to remove the selected object type. <p>If the objects have an input transformation applied, the transformed objects are the basis for the widget data.</p>
Metric	<p>Select a common metric or a metric for the selected object type in the list. The metric is the basis for the widget data.</p>

Option	Description
Label	<p>Type in a name that displays as a label for the metric.</p> <p>You can add a label if you have selected Metric Analysis > Top Highest Utilization or Metric Analysis > Top Lowest Utilization as Top-N options in the Configuration section.</p>
Unit	<p>You can define measurement units for the metrics. Select a measurement unit in the Unit drop-down menu.</p> <p>You can add a unit if you have selected Metric Analysis > Top Highest Utilization or Metric Analysis > Top Lowest Utilization as Top-N options in the Configuration section.</p>
Color Method	<p>You can use the Color Method option to define a coloring criteria for each metric. If this option is set to Custom, you can enter color values in the Yellow, Orange, and Red text boxes. If you do not want to use color, select None.</p> <p>You can add color thresholds if you have selected Metric Analysis > Top Highest Utilization, Metric Analysis > Top Lowest Utilization, or Metric Analysis > Percentile as Top-N options in the Configuration section.</p>
Output Filter	
Basic	<p>Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.</p> <p>If the objects have an input transformation applied, you select tag values for the transformed objects.</p>

Option	Description
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <p>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.
Additional Columns	<p>Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.</p> <ol style="list-style-type: none"> 1 Click the Add New Metrics icon to add metrics based on object types. The metrics that you add appear in a list in this section. <p>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the Select Object icon to select an object for the object type. Pick metrics of the selected object from the metric tree.</p> <p>For example, you can select the Datacenter object type, click the Select Object icon to display the list of data centers in your environment, and pick metrics of the selected data center.</p> <ol style="list-style-type: none"> 2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click Update.

Topology Graph Widget

The Topology Graph widget gives a graphical presentation of objects and their relationships in the inventory. You can customize each instance of the widget in your dashboard.



How the Topology Graph Widget and Configuration Options Work

The Topology Graph widget enables you to explore all nodes and paths connected to an object from your inventory. Connection between the objects might be a logical, physical, or network connection. The widget can display a graph that shows all of the nodes in the path between two objects, or that shows the objects related to a node in your inventory. You select the type of graph in the Exploration Mode when you configure the widget. You can select the levels of exploration between nodes in the displayed graph by using **Relationship** check boxes when you edit the widget. The widget displays all object types in the inventory by default, but you can select object types to view by using the Object View list during the configuration process. Double-clicking an object on the graph takes you to a detailed page about the object.

Where You Find the Topology Graph Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Topology Graph Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

Option	Description
Action	Use to select from predefined actions for each object type. To see available predefined actions, select an object in the graph and click the toolbar to select an action. For example, when you select a datastore object in the graph, you can click Delete Unused Snapshots for Datastore to apply this action to the object.
Dashboard Navigation	Takes you to a predefined object. For example, when you select a datastore from the graph and click Dashboard Navigation , you can open the datastore in the vSphere Web Client.
Pan	Use to move the entire graph.

Option	Description
Show values on point	Provides a tool tip with parameters when you point to an object in the graph.
Zoom in	Zooms in the graph.
Zoom out	Zooms out the graph.
Hierarchical View	Use to switch to hierarchical view. Hierarchical view is enabled only for Node Exploration mode and with selected inventory tree.
Graph View	Use to switch to graph view.
Object Detail	Select an object and click this icon to show the Object Detail page for the object.
Expand Node	Selects which object types related to your object to show on the graph. For example, if you select a virtual machine from the graph and click Expand Node toolbar icon and select Host System , the host on which the virtual machine is located is added to the graph.
Hide Node(s)	Use to remove a given object from the graph
Reset To Initial Object	Use to return to the initially displayed graph and configured object types.
Explore Node	Use to explore a node from a selected object in the graph. For example, if the graph displays a connection between a VM, a host, and a datastore, and you want to check the connection of the host with the other objects in the inventory, you can select the host and click Explore Node .
Status	Use to select objects based on their status or their state.

Topology Graph Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	

Option	Description
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	<p>If you enable the Refresh Content option, specify how often to refresh the data in this widget.</p>
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Exploration Mode	<p>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.</p> <p>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.</p> <p>Important To select object view is mandatory for the widget to start working in path exploration mode.</p>
Show Paths	<p>Use All to observe connections between a node and nodes related to it as well as connections between the nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph shows a VM connected to its datastore and host and the connection between the host and datastore.</p> <p>Use Discovered Only to observe directly related nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph will show the VM connected to its datastore and to its host, but without the connection between the host and datastore.</p>
Configuration File	<p>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.</p>
Metric Configuration	<p>Specifies a list with attributes to display.</p>

Option	Description
Layout	Select whether you want a graph view or hierarchical view for the topology graph.
Tree type	For a hierarchical layout, select whether you want a tree type view.
Input Data	
Selected object	From the object list, select an object on which you want to base the widget data.
Degree of separation	Available only when node exploration mode is selected. Use to define the levels of exploration in node exploration mode. The lowest degree configuration shows only directly related nodes rather than higher degrees that show the inventory in details.
Select First Object	Available only in path exploration mode. Select the first object from the object list.
Select Second Object	Available only in path exploration mode. Select the second object from the object list.
Object view	Use to select which types of objects to observe in the graph.
Relationship	Select the type of relationship between objects to observe in the graph, respectively the details about your inventory . The common relationships for all objects are parent and child, but the list of relationships can vary depending on added solutions to vRealize Operations Manager.

View Widget

The View widget provides the vRealize Operations Manager view functionality into your dashboard.

How the View Widget and Configuration Options Work

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, supermetrics, properties, alerts, policies, and data from a different perspective.

You can add the View widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. List views can send interactions to other widgets.

Where You Find the View Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

You can export the view as a CSV file for any view type.

View Widget Toolbar Options

The View widget toolbar depends on the displayed view type.

Option	Description
Export as CSV	You can export the view as a CSV file for any view type.
Open in External Application	Ability to link to another application for information about the object. For example, you have a List view with VMs. You can select any VM and select Open in External Application to open the VM in vSphere Web Client.
Time Settings	<p>Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.</p> <ul style="list-style-type: none"> ■ Relative Date Range. Select a relative date range of data transformation. ■ Specific Date Range. Select a specific date range of data transformation. ■ Absolute Date Range. Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month. <p>The units of time available are: Hours, Days, Weeks, Months, and Years.</p> <p>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.</p> <ul style="list-style-type: none"> ■ Dashboard Time. Select this option to enable the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.
Items per page	You can set the number of results that appear in the widget. Available for List view only.
Roll up interval	The time interval at which the data is rolled up.
Actions	An action on the selected object. Depends on the object type.
Filter	Limits the list to objects for a specific host, data center, and so on. You can drill-down in the hierarchical level. Available for List , Trend , and Distribution types of Views.
Filter by name	Limits the list to objects of a specific name. Available for List view only.

View Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Inventory trees	Select an existing predefined traversal spec to pick an object for the widget data.
Object	In self-provider mode, click the Add Object icon to select an object from the object list. The object list is displayed based on the inventory tree selection. You can also search for the object in this text box.
Output Data	
	<p>A list of defined views available for the selected object is displayed.</p> <p>You can create, edit, delete, clone, export, and import views directly from the View widget configuration options. For more information, see Views.</p>

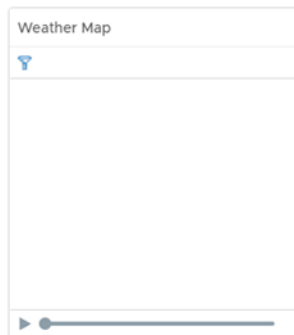
Option	Description
Auto Select First Row	Determines whether to start with the first row of data for list type views.
Show	<p>Select one or more of the following items to display in the widget:</p> <ul style="list-style-type: none"> ■ To display the list of legends in the widget, select Legend. ■ To display the name of the labels in the widget, select Labels.

Weather Map Widget

The Weather Map widget provides a graphical display of the changing values of a single metric for multiple resources over time. The widget uses colored icons to represent each value of the metric. Each icon location represents the metric value for particular resources. The color of an icon changes to show changes in the value of the metric.

How the Weather Map Widget and Configuration Options Work

You can add the Weather Map widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.



Watching how the map changes can help you understand how the performance of the metric varies over time for different resources. You can start or stop the display using the **Pause** and **Play** options at the bottom of the map. You can move the slider forwards or backwards to a specific frame in the map. If you leave the widget display and return, the slider remains in the same state.

The map does not show the real-time performance of the metrics. You select the time period, how fast the map refreshes, and the interval between readings. For example, you might have the widget play the metric values for the previous day, refreshing every half second, and have each change represent five minute's worth of metric values.

To view the object that an icon represents, click the object.

Where You Find the Weather Map Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Weather Map Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains the icons that you can use to view the graph.

Icon	Description
Pause and Play	Start or stop the display. The icon remains in the same state if you leave the widget display and return.
Display Filtering Criteria	View the current settings for the widget, including the current metric.

Weather Map Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.

Option	Description
Redraw Rate	<p>An interval at which cached data is refreshed based on newly collected data.</p> <p>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.</p> <p>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.</p>
Metric History	Select the time period for the weather map, from the previous hour to the last 30 days.
Metric Sample Increment	Select the interval between metric readings. For example, if you set this option to one minute and set the Metric History to one hour, the widget has a total of 60 readings for each metric.
Group by	Select a tag value by which to group the objects.
Sort by	Select Object name or Metric value to set the way to sort the objects.
Frame Transition Interval	Select how fast the icons change to show each new value. You can select the interval between frames and the number of frames per second (fps).
Start Over Delay	The number of seconds for the display to remain static when it reaches the end of the Metric History period, the most current readings, before it starts over again from the beginning.
Color	<p>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.</p> <p>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.</p> <p>If you set a minimum or maximum value, any metric at or beyond that value appears in the end color.</p>
Output Data	<p>Select an object type in your environment on which you want to base the widget data.</p> <ol style="list-style-type: none"> Click the Add Object Type icon to search for and add an object type. <p>When you search for object types, you can filter the types in the list by selecting a type from the Adapter Type drop-down menu or by using the Filter text box.</p> <ol style="list-style-type: none"> Optionally, select the object type from the list and click the Delete Object Type icon to remove the selected object type.

Option	Description
Metric	Select a common metric or a metric for the selected object type in the list. The metric will be the basis for the widget data. The object corresponding to the metric is the selected object for the widget.
Output Filter	
Basic	Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.
Advanced	<p>Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.</p> <p>If the objects have a tag filter applied in the Basic subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.</p> <ol style="list-style-type: none"> 1 In the first drop-down menu, select an object type. 2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select Metrics for the Datacenter object type, you can define a filter criteria based on the value of a specific metric for data centers. 3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects. 4 To add more filter criteria, click Add. 5 To add another filter criteria set, click Add another criteria set.

Workload Widget


The Workload widget displays data indicating how hard a selected resource is working.

The Workload widget displays a graph depicting how hard the object that you selected is working. The Workload widget reports data on CPU usage, Memory usage, Disk I/O, and Network I/O.

Where You Find the Workload Widget

The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.



About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.
- The metric named `devices|totalLatency_average` is not supported.

As a result, vRealize Operations Manager does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

Workload Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>

Option	Description
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

Workload Pattern Widget

The Workload Pattern widget displays a historical view of the hourly workload of an object.

Where You Find the Workload Pattern Widget


The widget might be included on any of your custom dashboards. In the menu, click **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, in the menu, click **Dashboards**. Click **Actions > Create Dashboard/Edit Dashboard** to add or edit a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

Workload Pattern

Workload Pattern

A historical view of hourly workload pattern of an object. This view helps you visualize if an object has been working hard over the last week and identify any hot spots which might cause performance issues.



Workload Pattern Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

Option	Description
Title	Enter a custom title that identifies this widget from other instances that are based on the same widget template.
Configuration	
Refresh Content	<p>Enable or disable the automatic refreshing of the data in this widget.</p> <p>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.</p>
Refresh Interval	If you enable the Refresh Content option, specify how often to refresh the data in this widget.
Self Provider	<p>Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.</p> <ul style="list-style-type: none"> ■ On. You define the objects for which data appears in the widget. ■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.
Input Data	
Object	Search for objects in your environment and select the object on which you are basing the widget data. You can also click the Add Object icon and select an object from the object list. You can use the Filter text box to refine the object list and the Tag Filter pane to select an object based on tag values.

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 4-160. Menu Options

Menu	Description
Dashboards	Lists the dashboard groups and the dashboards that are enabled. You can use this menu for a quick navigation through your dashboards. When you navigate to a dashboard using the Dashboards option, the dashboards are listed in the left pane of the Dashboards page. The dashboards are listed in the order in which they are selected, with the most recent dashboard selected, appearing at the top. You can reorder the dashboards on the left pane of the Dashboards page using drag and drop.
Shared	If the dashboard has been shared, the shared icon is displayed against the dashboard name.

Table 4-160. Menu Options (continued)

Menu	Description
Actions	<p>Available dashboard actions, such as edit, delete, remove dashboard from the menu, set as dashboard landing page, and set as the Home landing page. These actions are applied directly to the dashboard that you are on.</p> <p>To remove the dashboard as the Home landing page, from the dashboard that has been set as the Home landing page, select Actions > Reset from Home landing page.</p> <p>To remove the dashboard as the dashboard landing page, from the dashboard that has been set as the landing page, select Actions > Reset from Dashboards landing page.</p> <p>You can also create a dashboard and navigate to Manage Dashboards page from the Dashboards drop-down menu in the left pane.</p>
Dashboard Time	<p>The dashboard time panel is enabled by default on all predefined and user-created dashboards. Using this option, you can select a time for the widgets in the dashboard. The default time is 6 hours. The pre-defined time/day options in the panel are 1 hour, 6 hours, 24 hours, or 7 days. You can also set a customized time option.</p> <p>To enable widgets to use the dashboard time, select Date Controls/Time Range > Dashboard Time from the widget toolbar. Some widgets have Dashboard Time as the default option. For example, Metric Chart, View, Rolling View, Sparkline, Health Chart, and Mashup Chart widgets.</p> <p>Dashboard time persists if:</p> <ul style="list-style-type: none"> ■ You enable a widget in a dashboard to use the dashboard time and then log out and log back in, or ■ You enable a widget in a dashboard to use the dashboard time, and you export and then import the dashboard into another instance of vRealize Operations Manager.

Types of Dashboards

You can use the predefined dashboards or create your own custom dashboard in vRealize Operations Manager.

See [Chapter 5 Predefined Dashboards](#) for more information.

Custom Dashboards

You can create dashboards that meet your environment needs in vRealize Operations Manager.

For information about creating a dashboard, see [Create and Configure Dashboards](#).

Create and Configure Dashboards

To view the status of all objects in vRealize Operations Manager, create a dashboard by adding widgets or views. You can create and modify dashboards and configure them to meet your environment needs.

Procedure

- 1 In the menu, click **Dashboards**.
- 2 From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard**.
- 3 Complete the following steps to:
 - a Enter a name for the dashboard.
[Dashboard Name](#)
 - b Add widgets or views to the dashboard.
[Widget or View List Details](#)
 - c Configure widget interactions.
[Widget and View Interactions Details](#)
 - d Create dashboard navigation.
[Dashboard Navigation Details](#)
- 4 Click **Save**.
- 5 Click **Actions > Edit Dashboard** to modify the dashboard.

Dashboard Name

The name and visualization of the dashboard as it appears on the vRealize Operations Manager Home page.

Where You Add a Name in a Dashboard

To create your dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. Enter a name in the **New Dashboard** field.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

If you use a forward slash while entering a name, the forward slash acts as a group divider and creates a folder with the specified name in the dashboards list if the name does not exist. For example, if you name a dashboard **clusters/hosts**, the dashboard is named `hosts` under the group `clusters`.

Widget or View List Details

vRealize Operations Manager provides a list of widgets or views that you can add to your dashboard to monitor specific metrics and properties of objects in your environment.

Where You Add Widgets or Views to a Dashboard

To create your dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

How to Add Widgets or Views to a Dashboard

In the widgets list panel, you see a list of all the predefined vRealize Operations Manager widgets or views. Drag the widget or view to the dashboard workspace in the upper panel.

To locate a widget or view, you can enter the name or part of the name of a widget or view in the **Filter** option. For example, when you enter **top**, the list is filtered to display the Top Alerts, Top-N, and Topology Graph widgets. You can then select the widget you require.

Most widgets or views must be configured individually to display information. For more information about how to configure each widget, see [Widgets](#).

How to Arrange Widgets or Views in a Dashboard

You can modify your dashboard layout to suit your needs. By default, the first widgets or views that you add are automatically arranged horizontally wherever you place them.

- To position a widget or a view, drag the widget or view to the desired location in the layout. Other widgets and views automatically rearrange to make room.
- To resize a widget or a view, drag the bottom-right corner of the widget or the view.
- To maximize or minimize a widget or a view, use the maximize and minimize options in the top-right corner.

Widget and View Interactions Details

You can connect widgets and views so that the information they show depends on each other.

Where You Create Widget and View Interactions

To create interactions for widgets or views in a dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. From the toolbar, click **Show Interactions**.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.

How to Create and Remove Widget Interactions

The list of available interactions depends on the widgets or views in the dashboard. Widgets and views can provide, receive, and can both provide and receive interactions at the same time.

To create interactions, click **Show Interactions**. Click a provider plug and drag to the receiver. You can also apply interactions from receiver to provider plugs. For more information about how interactions work, see [Widget Interactions](#).

To remove interactions, click on the interaction line and select **Remove Interaction**. You can also click the provider plug and select **Remove Interaction > <widget name>**.

Dashboard Navigation Details

You can apply sections or context from one dashboard to another. You can connect widgets and views to widgets and views in the same dashboard or to other dashboards to investigate problems or better analyze the provided information.

Where You Add Another Dashboard

To create dashboard navigation to a dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Create Dashboard** to add a dashboard. In the dashboard workspace, click **Show Interactions**. From the **Select Another Dashboard** drop-down menu, select the dashboard to which you want to navigate.

To edit your dashboard, in the menu, click **Dashboards**. Click **Actions > Edit Dashboard** to edit the selected dashboard.


How Dashboard Navigation Works

You can create dashboard navigation only for provider widgets and views. The provider widget or view sends information to the destination widget or view. When you create dashboard navigation, the destination widgets or views are filtered based on the information type they can receive.

How to Add Dashboard Navigation to a Dashboard

The list of available dashboards for navigation depends on the available dashboards and the widgets and views in the current dashboard. To add navigation, you can drag from a sender widget interaction plug to a receiver widget interaction plug. You can select more than one applicable widget or view.

Note If a dashboard is unavailable for selection, it is unavailable for dashboard navigation.

The Dashboard Navigation icon () appears in the top menu of each widget or view when a dashboard navigation is available.

After you have set widget interaction in the provider dashboard, the widget and menu bar are highlighted and two arrows appear in the top-left corner of the widget. After you have set widget interaction, clicking the object in the provider widget takes you to the receiver widget of the navigated dashboard.

Manage Dashboards

You can select dashboards individually or as a group and perform several actions.

To manage your dashboards, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Use the options from the horizontal ellipsis next to the **Add** option.

All the dashboards are listed on this page. You can filter the dashboards based on the name of the dashboard, the dashboard folder, enabled dashboards, shared dashboards, or the dashboard owner. You can click **Add** to create a dashboard. For information about creating a dashboard, see [Create and Configure Dashboards](#).

You can select a dashboard from the list, click the vertical ellipsis against each dashboard, and select the various options such as edit, delete, clone, and disable a dashboard. You can also change ownership of dashboards, save the dashboard as a template, and export the dashboard. By default, the list of dashboards is sorted by name and all the columns can be sorted.

Note A wrench icon appears when the data in an imported dashboard depends on the existence of one or more adapters that are currently not present. The wrench icon disappears if the required data in an imported dashboard appears in vRealize Operations Manager after configuration.

Imported dashboards regardless of used data, remain stuck and include a wrench icon if the dashboard that is stuck (with the wrench icon), already exists.

Datagrid Options

Column Names	Description
Name	Displays the name of the dashboard.
Folder	Lists the folder to which each dashboard belongs.
Description	Displays the description of the dashboard.
Enabled	Enables and disables the dashboard.
URL	Displays whether the dashboard is shared externally. For dashboards that have been shared, click to view the shared links.
Shared	Displays whether the dashboard is shared internally. Click to view and edit the groups to which the dashboard has been shared.
Owner	Displays the owner of the dashboard.
Last Modified	Displays the date the dashboard was last modified.

You can select more than one dashboard and perform a set of options by clicking the horizontal ellipsis next to the **Add** option.

Table 4-161. Dashboards Options

Option	Description	Usage
Export	When you export a dashboard, vRealize Operations Manager creates a dashboard file in JSON format.	You can export a dashboard from one vRealize Operations Manager instance and import it to another.
Enable	Enables a dashboard that was previously disabled.	

Table 4-161. Dashboards Options (continued)

Option	Description	Usage
Disable	Disables a dashboard.	
Delete	Deletes a dashboard.	
Change Ownership	Assigns a new owner to the dashboard.	After you assign a dashboard to a new owner, the dashboard is no longer displayed as one of your dashboards. When you transfer a dashboard that was previously shared with user groups, information about the shared user groups and group hierarchy is retained.
Import	A PAK or JSON file that contains dashboard information from vRealize Operations Manager.	You can import a dashboard that was exported from another vRealize Operations Manager instance.
Auto-rotate Dashboards	Changes the order of the dashboard tabs on vRealize Operations Manager home page.	You can configure vRealize Operations Manager to switch from one dashboard to another. For more information, see Auto-Rotate Dashboards .
Manage Summary Dashboards	Provides you with an overview of the state of the selected object, group, or application.	You can change the Summary tab with a dashboard to get information specific to your needs. For more information, see Manage Summary Dashboards
Manage Dashboard Folders	Groups dashboards in folders.	You can create dashboard folders to group the dashboards in a way that is meaningful to you. For more information, see Manage Dashboard Folders .
Manage Dashboard Sharing	Makes a dashboard available to other users or user groups.	You can share a dashboard or dashboard template with one or more user groups. For more information, see Share Dashboards with Users .
Clone	Copies a dashboard to other another user or user group.	You can copy a dashboard to another user or user group. Specify the dashboards to be shared and select a target user and specify the target folder. Accessible as an option only from the vertical ellipsis against the selected dashboard.
Save as Template	Contains all the information in a dashboard definition.	You can use any dashboard to create a template. Accessible as an option only from the vertical ellipsis against the selected dashboard.

The dashboard list depends on your access rights.

Manage Summary Dashboards

The **Summary** tab provides you with an overview of the state of the selected object, group, or application. You can change the **Summary** tab with a dashboard to get information specific to your needs.

Where You Configure a Summary Tab Dashboard

To manage the summary dashboards, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and select **Manage Summary Dashboards**.

How You Manage the Summary Dashboards

Table 4-162. Manage Summary Dashboards Toolbar Options

Option	Description
Use Default	Click to use vRealize Operations Manager default Summary tab.
Assign a Dashboard	Click to view the Dashboard List dialog box that lists all the available dashboards.
Adapter Type	Adapter type for which you configure a summary dashboard.
Filter	Use a word search to limit the number of adapter types that appear in the list.

To change the Summary tab for an object, select the object in the left panel, click the **Assign a Dashboard** icon. Select a dashboard for it from the All Dashboards dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard that you have associated to the object type when you navigate to the **Summary** tab of the object details page.

Auto-Rotate Dashboards

You can change the order of the dashboard tabs on your home page. You can configure vRealize Operations Manager to switch from one dashboard to another. This feature is useful if you have several dashboards that show different aspects of your enterprise's performance and you want to look at each dashboard in turn.

Where You Configure Auto-Rotation of a Dashboard

To reorder and configure a dashboard switch, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Select **Auto-rotate Dashboards** from the horizontal ellipsis next to the **Add** option.

How You Reorder the Dashboards

The list shows the dashboards as they are ordered. Drag the dashboards up and down to change their order on the home page.

How You Configure an Automatic Dashboard Rotation

- 1 Double-click a dashboard from the list to configure.
- 2 From the Rotation drop-down menus, select **On**.
- 3 Select the time interval in seconds.
- 4 Select the dashboard to switch and click **Update**.
- 5 Click **Save** to save your changes.

On the home page, the current dashboard will switch to the dashboard that is defined after the specified time interval.

Manage Dashboard Folders

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

Where You Manage Dashboard Folders

To manage the dashboard folders, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Folders**.

How You Manage the Dashboard Folders

Table 4-163. Manage Dashboard Folders Options

Option	Description
Dashboards List	A list with all available dashboards.
Folders	A hierarchy tree with all the available group folders.

To create a dashboard folder, click **New Folder** in the **Folders** pane and enter the name of the folder. If you want to create a folder under another folder, select a parent folder under which you want to create the child folder, then click **New Folder**. To add a dashboard, drag one from the dashboards list to the selected folder in the **Folders** pane.

You can delete folders and/or detach dashboards from a folder, by selecting one or more folders and dashboards from the **Folders** pane and by clicking **Actions > Delete**.

You can rename a folder by selecting a single folder from the **Folders** pane and by clicking **Actions > Rename**.

Share Dashboards with Users

You can share a dashboard with one or more user groups. When you share a dashboard, it becomes available to all the users in the user group that you select. The dashboard appears the same to all the users who share it. If you edit a shared dashboard, the dashboard changes for all users. Other users can only view a shared dashboard. They cannot change it.

Where You Share a Dashboard From

To share a dashboard, in the menu, click **Dashboards**. From the left pane, click the **Dashboards** drop-down menu and then click **Manage Dashboards**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Sharing**.

Table 4-164. Dashboard Sharing Options

Option	Description
All Dashboards	Link to view all the available dashboards that you can share. The dashboards are displayed on the right side in the dashboards list.
User Groups	Lists the available user groups that you can share a dashboard with. The list includes the Everyone group.
Dashboard List	List of shared dashboards with the selected user group or all the available dashboards that you can share, if no user group is selected.

Manage Dashboard Sharing

To share a dashboard, navigate to the dashboard in the list of dashboards and drag it to the group to share it with, on the left.

To stop sharing a dashboard with a group, click that group on the left panel, navigate to the dashboard in the right panel, and click **Stop Sharing** above the list.

Dashboards Actions and Options

You can change the order of the dashboard tabs, configure vRealize Operations Manager to switch from one dashboard to another, create dashboard folders to group the dashboards in a way that is meaningful to you, share a dashboard or dashboard template with one or more user groups, and transfer selected dashboards to a new owner.

Options for Sharing Dashboards

You can share predefined or custom dashboards using URLs, emails, and by copying the code to embed the dashboard into confluence or other internal official web pages. You can also assign and unassign a dashboard to specific user groups and export the dashboard configuration details.

When you use a non-authenticated shared URL, as a user you can open the dashboard in a new browser session. If you have already logged into vRealize Operations Manager in another session, you are redirected to this dashboard and the user authentication permissions apply. To ensure that the non-authenticated URL opens the intended dashboard, as a user you must log out from all existing user sessions.

The dashboard shared with the URL opens in a page where you can access all the widgets within the dashboard and you can interact with the given widgets at the same time. A non-authenticated dashboard however, does not allow you to browse to other areas of vRealize Operations Manager.

Dashboard sharing can only be applied to Groups with a vRealize Operations Standard Edition license.

Where You Can Access the Options to Share Dashboards

From the menu, select **Dashboards**. Click on an existing dashboard and then click the **Share Dashboard** icon in the top-right corner.

Table 4-165. Options in the Share Dashboard Dialog Box

Option	Description
URL	<p>Allows you to copy the tiny URL for the selected dashboard.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 Months, or Never Expire. ■ Click Copy Link to copy the link to a new window from where you can view the dashboard. <hr/> <p>Note</p> <ul style="list-style-type: none"> ■ As a user, if you open a shared link and you are logged into vRealize Operations Manager, you are navigated to your default dashboard, instead of viewing the shared one. ■ As a user, if you log in to the same IP that was shared with you previously, you cannot access the page with the same browser. ■ As a user, ensure that you have the following permission: Dashboards > Dashboard Management > Share (Public). <hr/> <p>You can stop sharing a dashboard you had previously shared. To stop sharing a dashboard, click the Unshare Link option and enter the URL of the dashboard that you want to stop sharing and click Unshare.</p> <p>Authentication is not required to view the shared dashboard.</p>
Email	<p>Allows you to send an email with the URL details of the dashboard, to a specific person.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 months, or Never Expire. ■ Configure an SMTP instance. See Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts. ■ Enter an email address and click the Send Email button to send an email with the URL details of the dashboard. <p>Authentication is not required to view the shared dashboard.</p>

Table 4-165. Options in the Share Dashboard Dialog Box (continued)

Option	Description
Embed	<p>Provides an embedded code for the dashboard. You can use this code to embed the dashboard in relevant confluence pages that your company executives routinely use and analyze.</p> <ul style="list-style-type: none"> ■ Set the expiry period for the link to 1 day, 1 week, 1 month, 3 Months, or Never Expire. <hr/> <p>Note</p> <ul style="list-style-type: none"> ■ If you embed a dashboard in the Text widget, the widget does not display any data. ■ When you open an HTML/confluence page with an embedded dashboard from the same browser that you have logged into vRealize Operations Manager, the dashboard does not load. <hr/> <p>Authentication is not required to view the shared dashboard.</p>
Groups	<p>Allows you to assign and unassign a dashboard to specific user groups.</p> <ul style="list-style-type: none"> ■ Select the group to which you want to grant dashboard access from the drop-down menu and click Include. You can include more than one dashboard. ■ From the label, select the cross mark to unassign the dashboard. <p>Log in to vRealize Operations Manager to view the shared dashboard.</p>
Export	<p>Allows you to export the dashboard configuration details.</p> <p>Log in to vRealize Operations Manager to export/import a dashboard.</p>

Manage Widgets in Dashboards

You can replicate widgets multiple times in a dashboard by using the copy and paste functionality.

Navigate to the dashboard from which you want to copy widgets. Select **Actions > Edit Dashboard**. Select one or more widgets that you want to copy by clicking the title of the widget and then select **Actions > Copy Widget(s)**. Click **Actions > Paste Widget(s)** to paste one or more widgets in the same dashboard.

To paste one or more widgets into another dashboard, exit the edit screen of the dashboard by selecting **Cancel**. Navigate to the dashboard to which you want to paste one or more widgets and select **Actions > Edit Dashboard** and then **Actions > Paste Widget(s)**.

Views

vRealize Operations Manager provides several types of views. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms,

and so on, from a different perspective. Views also show information that the adapters in your environment provide.

You can configure vRealize Operations Manager views to show transformation, forecast, and trend calculations.

- The transformation type determines how the values are aggregated.
- The trend option shows how the values tend to change, based on the historical, raw data. The trend calculations depend on the transformation type and roll up interval.
- The forecast option shows what the future values can be, based on the trend calculations of the historical data.

You can use vRealize Operations Manager views in different areas of vRealize Operations Manager.

- To manage all views, in the menu, click **Dashboards**, and then in the left pane click **Views**.
- To see the data that a view provides for a specific object, navigate to that object, click the **Details** tab, and click **Views**.
- To see the data that a view provides in your dashboard, add the View widget to the dashboard. For more information, see [View Widget](#).
- To have a link to a view in the Further Analysis section, select the Further Analysis option on the view workspace visibility step.

Views and Reports Ownership

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Views Overview

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective.

How You Access the Views Page

In the menu, click **Dashboards**, and then in the left pane click **Views** to access the Views page.

Manage and Preview Views

You can preview a view by clicking a view from the **Views** page. Add an object if necessary, by clicking **Select preview source** from the upper-right corner of the **Views** page. The preview of the view appears just below the **Views** option in the right pane.

You can select a view from the list, click the vertical ellipsis against each view, and select the various options such as edit, delete, clone, and export a view.

You can filter the views based on the name, type, description, subject, and owner. You can click **Add** to create a view. For information about creating a view, see [Create and Configure a View](#).

You can select more than one view and delete, export, and import views by clicking the horizontal ellipsis next to the **Add** option.

Views are also categorized and listed in the **Views** menu based on the type of view and subject. You can access the **Views** menu from a specific view preview page.

Table 4-166. Filter Groups

Filter Group	Description
Name	Filter by the view name. For example, type my view to list all views that contain the my view phrase in their name.
Type	Filter by the view type.
Description	Filter by the view description. For example, type my view to list all views that contain the my view phrase in their description.
Subject	Filter by the subject.
Owner	Filter by owner.

Views and Reports Ownership

The owner of views, reports, or templates might change over time.

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Create and Configure a View

To collect and display information for a specific object, you can create a custom view.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Click **Add** to create a view.
- 3 Complete the steps in the left pane to:
 - a Enter a name and description for the view.
[Name and Description Details](#)
 - b Change the presentation of a view.
[Presentation Details](#)
 - c Select the base object type for a view.
[Subjects Details](#)
 - d Add data to a view.
[Data Details](#)
 - e Change the visibility of a view.
[Visibility Details](#)
- 4 Click **Save**.

Name and Description Details

The name and description of the view as they appear in the list of views on the Views page.

To add a name and description to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. Click the **Add** option. In the workspace, on the left, click **Name and Description**.

Table 4-167. Name and Description Options in the View Workspace

Option	Description
Name	Name of the view as it appears on the Views page.
Description	Description of the view.

Presentation Details

A presentation is a way the collected information for the object is presented. Each type of view helps you to interpret metrics and properties from a different perspective.

To change the presentation of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. Click the **Add** option. In the workspace, on the left, click **Presentation**. If you create a view, complete the required previous steps.

Table 4-168. Presentation Options in the View Workspace

View Type	Description
List	Provides tabular data about specific objects in the monitored environment. Column count is limited to 25 in a PDF report and 50 in a CSV report. Page count is unlimited.
Summary	Provides tabular data about the use of resources in the monitored environment.
Trend	Uses historic data to generate trends and forecasts for resource use and availability in the monitored environment.
Distribution	Provides aggregated data about resource distribution in the monitored environment. When you add a distribution type of View to a dashboard, you can click a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment.
Text	Inserts the provided text. The text can be dynamic and contain metrics and properties. You can format text to increase or decrease the font size, change the font color, highlight text, and align text to the left, right, or center. You can also make the selected text appear bold, in italics, or underlined. By default the text view is available only for report template creation and modification. You can change this on the Visibility step of the view workspace.
Image	Inserts a static image. By default the image view is available only for report template creation and modification. You can change this on the Visibility step of the view workspace.

You can see a live preview of the view type when you select a subject and data, and **Select preview source**.

How to Configure the Presentation of a View

Some of the view presentations have specific configuration settings.

Table 4-169. Presentation Configuration Options in the View Workspace

View Type	Configuration Description
List	<ul style="list-style-type: none"> ■ Select the number of items per page. Each item is one row and its metrics and properties are the columns. ■ Select the top results. Restricts the number of results. For example, if you list all the clusters in a View, selecting 10 in this option displays the top 10 clusters with the relevant information. You can reduce the number of rows for the purposes of reporting.
Summary	Select the number of items per page. Each row is an aggregated metric or property.
Trend	<p>Enter the maximum number of plot lines. Limits the output in terms of the objects displayed in the live preview of the view type on the left upper pane. The number you set as the maximum number of plot lines determines the plot lines.</p> <p>For example, if you plot historical data and set the maximum at 30 plot lines, then 30 objects are displayed. If you plot historical, trend, and forecast lines, and set the maximum to 30 plot lines, then only 10 objects are displayed as each object has three plot lines.</p>
Distribution	<p>Select the visualization of the distribution information in a pie chart or a bar chart.</p> <p>Select the distribution type, and configure the buckets count and size.</p> <p>To understand vRealize Operations Manager distribution type, see View Distribution Type.</p>

Coloring

Configuration Option	Description
Colorize	The colors of the slices in the pie chart are displayed in the order of the colors in the color palette.
Select Color	Select the color that you want the chart to appear in. If there is more than one slice in a pie chart, the colors are chosen sequentially from the color palette. In a bar chart, the bars are all the same color.

Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Dynamic distribution

You specify in details how vRealize Operations Manager distributes the data in the buckets.

Table 4-170. Dynamic Distribution Configuration Options

Configuration Option	Description
Buckets Count	The number of buckets to use in the data distribution.
Buckets Size Interval	The bucket size is determined by the defined interval divided by the specified number of buckets.
Buckets Size Logarithmic bucketing	The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data.
Buckets Size Simple Max/Min bucketing	The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets.

Manual distribution

You specify the number of buckets and the minimum and maximum values of each bucket.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distribute the data.

View Distribution Type

vRealize Operations Manager view distribution type provides aggregated data about resource distribution in the monitored environment.

Visualization

You can view the data as a pie chart, a bar chart, or a donut chart. When you add a distribution type of View to a dashboard, you can click a section of the pie chart, or on one of the bars in the bar chart, or a section of the donut chart to view the list of objects filtered by the selected segment. You can select the display colors for single or multi-colored charts.

Dynamic distribution

You specify in details how vRealize Operations Manager distributes the data in the buckets.

Table 4-171. 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 4-171. 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. You can also select a color for each defined bucket that you specify.

Discrete distribution

You specify the number of buckets in which vRealize Operations Manager distributes the data. If you increase the number of buckets, you can see more detailed data.

Subjects Details

The subject is the base object type for which the view shows information.

To specify a subject for a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. Click the **Add** option. In the workspace, on the left, click **Subjects**. If you create a view, complete the required previous steps.

The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. You can limit the level where the view appears with the Blacklist option in the **Visibility** step.

View availability depends on the view configuration subject, inventory view, user permissions, and view Visibility settings.

For list views with **Symptom** as a subject, the following columns can be sorted: Criticality Level, Status, Object Type, Object Name, Created on, and Canceled on. You cannot sort the Triggered On and Violation Info columns. If other symptom metrics exist, you cannot sort any of the columns.

In a List view, you can group the results based on a parent object, by making a selection in the **Group By** drop-down option. If you generate a report based on the list view for which a group has been specified, the report displays group-based information for the selected object. You can also view summary calculations for the group of objects in the report, along with the total summary results for all the objects.

Views Applicability

Views might not always appear where you expect them to. The main applicability of views depends on the view subject and the inventory view.

List View

When you navigate through the environment tree, you can see the List view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the List view might be missing at the object containers. For example, you create a List view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Storage > vSphere World**, select the same vCenter Server, and click the **Details** tab, your List view is missing. Your List view with subject Host System is missing because the object Host System is not included in the vSphere Storage inventory view.

Summary View

When you navigate through the environment tree, you can see the Summary view at the subjects that you specify during the view configuration and at their object containers. Depending on the inventory view, the Summary view might be missing at the object containers. For example, you create a Summary view with subject Datastore. When you go to **Environment > vSphere Storage > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your List view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Summary view is missing. Your Summary view with subject datastore is missing because the object Datastore is not included in the vSphere Networking inventory view.

Trend View

When you navigate through the environment tree, you can see the Trend view only at the subjects that you specify during the view configuration. For example, you create a Trend view with subject Virtual Machine. When you navigate to a virtual machine in the navigation tree, you see your view.

Distribution View

When you navigate through the environment tree, you can see the Distribution view only at the object containers of the subjects that you specify during the view configuration. Depending on the inventory view, the Distribution view might be missing at the object containers. For example, you create a Distribution view with subject Host System. When you go to **Environment > vSphere Hosts and Clusters > vSphere World**, select a vCenter Server, and click the **Details** tab, you can see your Distribution view. If you go to **Environment > vSphere Networking > vSphere World**, select the same vCenter Server, and click the **Details** tab, your Distribution view is missing. Your Distribution view with subject Host System is missing because the object Host System is not included in the vSphere Networking inventory view.

Text View

When you navigate through the environment tree, you can see the Text view only at the subjects that you specify during the view configuration. For example, you create a Text view with subject vCenter Server. When you navigate to a vCenter Server in the navigation tree, you see your view. If you did not specify a subject, you see your view for every subject in the environment.

Image View

The Image view is applicable for every object in the environment.

Note Views applicability depends also on your user permissions and the view Visibility configuration.

Data Details

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which vRealize Operations Manager collects, calculates, and presents the information for the view.

To add data to a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. Click the **Add** option. In the workspace, on the left, click **Data**. If you create a view, complete the required previous steps.

How to Add Data to a View

If you selected more than one subject, specify the subject for which you add data. Double-click the data from the tree in the left panel to add it to the view. For each subject the data available to add might be different.

How to Configure the Data Transformation

The data configuration options depend on the view and data type that you select. Most of the options are available for all views.

Table 4-172. 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 4-172. Data Configuration Options (continued)

Configuration Option	Description
Transformation	<p>Determines what calculation method is applied on the raw data. You can select the type of transformation:</p> <ul style="list-style-type: none"> ■ Minimum. The minimum value of the metric over the selected time range. ■ Maximum. The maximum value of the metric over the selected time range. ■ Average. The mean of all the metric values over the selected time range. ■ Sum. The sum of the metric values over the selected time range. ■ First. The first metric value for the selected time range. ■ Last. The last value of a metric within the selected time range. <p>If you have selected Last as the transformation in versions before vRealize Operations Manager 6.7, and the end of specified time range is not before the last five minutes, use the Current transformation.</p> <ul style="list-style-type: none"> ■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null. ■ Standard Deviation. The standard deviation of the metric values. ■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum. <p>The time period accuracy (based on which the nearest point to the extremum is taken in the original data) is calculated over the correlated metric time stamps. In an ideal case, it represents half of the collection cycle of the correlated metric, for example Tacc. The Metric Correlation transformation takes the time stamp of the extremum point in the correlated metric data, for example T, and then defines the following time range: [T - Tacc, T + Tacc]. It then looks for any value within that range in the original metric data, and if not found, returns null.</p> <ul style="list-style-type: none"> ■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range. ■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on. ■ Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, <code>sum / ((max + min) / 2)</code>. You can use the operands of

Table 4-172. Data Configuration Options (continued)

Configuration Option	Description
	<p>some of the existing transformations such as, max, min, avg, sum, first, last, current. You cannot use standard deviation, forecast, metric correlation, and percentile.</p> <p>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps.</p> <p>Available for all views, except Trend.</p> <ul style="list-style-type: none"> ■ Timestamp: You can choose between Absolute Timestamp OR Relative Timestamp. ■ If applied to a numeric metric/property defined with a time-unit definition, the actual value is converted to a human readable timestamp. The metric value is rounded-off to an hour. <p>Applicable for Absolute Timestamp.</p> <ul style="list-style-type: none"> ■ In the remaining cases, a timestamp is displayed when metrics and properties are added or modified. In this case, the behavior is the same as the Timestamp option selected for a non-Timestamp transformation. <p>Applicable for Absolute Timestamp and Relative Timestamp.</p> <p>Available for List view and Minimum, Maximum, Current, First, and Last transformations.</p>
Ranges for metric coloring	<p>You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the Red Bound field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats.</p>
Data Series	<p>You can select whether to include historical data, trend of historical data, and forecast for future time in the trend view calculations.</p> <p>Available for Trend view.</p>

Table 4-172. Data Configuration Options (continued)

Configuration Option	Description
Series Roll up	<p>The time interval at which the data is rolled up. You can select one of the available options. For example, if you select <code>Sum</code> as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.</p> <p>This option is applicable to the Transformation configuration option.</p> <p>Available for all views.</p>
Threshold Lines	<p>You can set a threshold for a single metric:</p> <ul style="list-style-type: none"> ■ None. You have not set a threshold. ■ By Symptom Definition. You can set a threshold value based on a symptom definition. ■ Custom. You can set the threshold value as Warning, Critical, or Immediate. These options are available only for the Custom option. <p>Available for Trend view.</p>

How to Configure Time Settings

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data.

Table 4-173. Time Settings Options

Configuration Option	Description
Time Range Mode	<p>In Basic mode, you can select date ranges.</p> <p>In Advanced mode, you can select any combination of relative or specific start and end dates.</p>
Relative Date Range	<p>Select a relative date range of data transformation.</p> <p>Available in Basic mode.</p>
Specific Date Range	<p>Select a specific date range of data transformation.</p> <p>Available in Basic mode.</p>

Table 4-173. Time Settings Options (continued)

Configuration Option	Description
Absolute Date Range	<p>Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.</p> <p>The units of time available are: Hours, Days, Weeks, Months, and Years.</p> <p>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.</p> <p>Available in Basic mode.</p>
Relative Start Date	<p>Select a relative start date of data transformation.</p> <p>Available in Advanced mode.</p>
Relative End Date	<p>Select a relative end date of data transformation.</p> <p>Available in Advanced mode.</p>
Specific Start Date	<p>Select a specific start date of data transformation.</p> <p>Available in Advanced mode.</p>
Specific End Date	<p>Select a specific end date of data transformation.</p> <p>Available in Advanced mode.</p>
Currently selected date range	<p>Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM.</p>

How to Break Down Data

You can break down data in List views by adding interval or instance breakdown columns from the **Group By** tab.

Table 4-174. Group By Options

Option	Description
Add interval breakdown column (see data for column settings)	<p>Select this option to see the data for the selected resources broken down in time intervals.</p> <p>In the Data tab, select Interval Breakdown to configure the column. You can enter a label and select a breakdown interval for the time range.</p>
Add instance breakdown column (see data for column settings)	<p>Select this option to see the data for all instances of the selected resources.</p> <p>In the Data tab, select Instance Name to configure the column. You can enter a label and select a metric group to break down all the instances in that group. Deselect Show non-instance aggregate metric to display only the separate instances. Deselect Show only instance name to display the metric group name and instance name in the instance breakdown column.</p> <p>For example, you can create a view to display CPU usage by selecting the metric CPU:0 Usage. If you add an instance breakdown column, the column CPU:0 Usage displays the usage of all CPU instances on separate rows (0, 1, and so on). To avoid ambiguity, you can change the metric label of CPU:0 Usage to Usage.</p>

How to Add a Filter

The filter option allows you to add additional criteria when the view displays too much information. For example, a list view shows information about the health of virtual machines. From the **Filter** tab, you add a risk metric less than 50%. Then the view shows the health of all virtual machines with risk less than 50%.

To add filter to a view, select **Content > Views** in the left pane. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data** and click the **Filter** tab in the main panel. If you create a view, complete the required previous steps.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 4-175. Filter Add Options

Option	Description
Add	<p>Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.</p> <p>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.</p> <p>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum.</p>
Add another criteria	Adds another criteria set. The filter returns results that match one criteria set or another.

How to Add a Summary Row or Column to a View

The summary option is available only for List and Summary views. It is mandatory for the Summary views. You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, select **Content > Views** in the left pane. On the Views toolbar, click the plus sign to add a view or the pencil to edit the selected view. In the workspace, on the left, click **Data** and click the **Summary** tab in the main panel. If you create a view, complete the required previous steps.

For the List view, the summary row shows aggregated information by the specified subjects.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

Visibility Details

The view visibility defines where you can see a view in vRealize Operations Manager.

To change the visibility of a view, in the menu, click **Dashboards**, and then in the left pane click **Views**. Click the **Add** option. In the workspace, on the left, click **Visibility**. If you create a view, complete the required previous steps.

Table 4-176. View Workspace Visibility Options

Option	Description
Availability	Select where in vRealize Operations Manager you want to see this view. If you want to have the view available in a dashboard, select the check box, add the View widget, and configure it. You can also make the view available in report templates and in the Detail tab of a specific object when you select the specific check box.
Further Analysis	Select the Compliance check box to make the view available in the Compliance tab for a specific object.
Blacklist	Select a subject level where you do not want to see this view. For example, you have a list view with subject virtual machines. It is visible when you select any of its parent objects. You add data center in the banned list. The view is not visible anymore on data center level.

Editing, Cloning, and Deleting a View

You can edit, clone, and delete a view. Before you do, familiarize yourself with the consequences of these actions.

Edit a View

When you edit a view, all changes are applied to the report templates that contain it. To edit a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Edit**.

Clone a View

When you clone a view, the changes that you make to the clone do not affect the source view. To clone a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Clone**.

Delete a View

When you delete a view, it is removed from all the report templates that contain it. To delete a view, from the main menu, click **Dashboards**, and then from the left panel click **Views**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Delete**.

Including Deleted VMs in List View

In vRealize Operations Manager, you can view the deleted objects and the relationship of the objects in the list view. The objects can be VMs, deployments, projects, vApps, and edge gateways. You can also retain the relationship of the objects even after the objects are deleted from the system. The cost of the deleted virtual machines (VMs) is available until the retention period for that VM is over.

Where You Find Global Settings for Deleted VMs

To specify for how long you want to retain the deleted virtual machines in vRealize Operations Manager, navigate to **Administration > Management > Global Settings > Deleted Objects**.

You can also specify the **Deletion Scheduling Interval** which specifies the number of hours between resource deletion scheduling.

In the **Object Deletion Scheduling** page, click **Add**, select the virtual machine object from the drop-down menu, specify the value, and click **Update**. The global setting value for the deleted virtual machine is updated in vRealize Operations Manager.

For vRealize Automation, the price of the deleted VMs or deployments is added to the corresponding project object as a separate metric. If the deleted VM from vRealize Automation is associated with a cost-based pricing policy, then the price for that VM is not added to the corresponding project.

For vCloud Director, the price of deleted VMs, vApps, and Edge Gateways is added to the corresponding organization VDC object again as a separate metric. For vCenter Server, if VM is on unclustered Host, then deleted VM price is assigned to the Host, otherwise to the Cluster.

How to Include Deleted VMs in List View

To view the deleted virtual machines in the list view, navigate to **Dashboards > Views > Deleted VMs**. Click **Actions** and select **Edit View**. From the filter tab, select **Included Deleted Objects**.

The deleted VMs are visible in **Administration > Inventory > Object Types > Collection States > Not Existing**.

User Scenario: Create, Run, Export, and Import a vRealize Operations Manager View for Tracking Virtual Machines

As a virtual infrastructure administrator, you use vRealize Operations Manager to monitor several environments. You must know the number of virtual machines on each vCenter Server instance. You define a view to gather the information in a specific order and use it on all vRealize Operations Manager environments.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

You will create a distribution view and run it on the main vRealize Operations Manager environment. You will export the view and import it in another vRealize Operations Manager instance.

Procedure

1 Create a vRealize Operations Manager View for Supervising Virtual Machines

To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

2 Run a View

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

3 Export a View

To use a view in another vRealize Operations Manager instance, you export a content definition XML file.

4 Import a View

To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Create a vRealize Operations Manager View for Supervising Virtual Machines

To collect and display data about the number of virtual machines on a vCenter Server, you create a custom view.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.

- 2 Click the **Add** option to create a view.

- 3 Enter **Virtual Machines Distribution**, the name for the view.

- 4 Enter a meaningful description for the view.

For example, **A view showing the distribution of virtual machines per hosts.**

- 5 Click **Presentation** and select the **Distribution** view type.

The view type is the way the information is displayed.

- a From the **Visualization** drop-down menu, select **Pie Chart**.

- b From the Distribution Type configurations, select **Discrete distribution**.

Leave **Max number of buckets** deselected because you do not know the number of hosts on each vCenter Server instance. If you specify a number of buckets and the hosts are more than that number, one of the slices shows unspecified information labeled Others.

- 6 Click **Subjects** to select the object type that applies to the view.

- a From the drop-down menu, select **Host System**.

The Distribution view is visible at the object containers of the subjects that you specify during the view configuration.

- 7 Click **Data** and in the filter text box enter **Total Number of VMs**.

- 8 Select **Summary > Total Number of VMs** and double-click to add the metric.

- 9 Retain the default metric configurations and click **Save**.

Run a View

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Environment**.
- 2 In the left pane, navigate to a vCenter Server instance and click the **Details** tab.
All listed views are applicable for the vCenter Server instance.
- 3 From the **All Filters** drop-down menu on the left, select **Type > Distribution**.
You filter the views list to show only distribution type views.
- 4 Navigate to and click the **Virtual Machines Distribution** view.
The bottom pane shows the distribution view with information about this vCenter Server. Each slice represents a host and the numbers on the far left show the number of virtual machines.

Export a View

To use a view in another vRealize Operations Manager instance, you export a content definition XML file.

If the exported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Select a view and click **Export** from the vertical ellipsis next to the selected view.

Import a View

To use views from other vRealize Operations Manager environments, you import a content definition XML file.

Prerequisites

Verify that you have the necessary access rights to perform this task. Your vRealize Operations Manager administrator can tell you which actions you can perform.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Views**.
- 2 Select a view and click the **Import** option from the horizontal ellipsis.
- 3 Browse to select the Virtual Machines Distribution content definition XML file and click **Import**.

If the imported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

Note The imported view overwrites if a view with the same name exists. All report templates that use the existing view are updated with the imported view.

Reports

A report is a scheduled snapshot of views and dashboards. You can create it to represent objects and metrics. It can contain table of contents, cover page, and footer.

With the vRealize Operations Manager reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.

Report Templates Tab

On the **Report Templates** tab you can create, edit, delete, clone, run, schedule, export, and import templates.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Report Templates** to access the Reports Templates tab.

All templates that are applicable for the selected object are listed on the **Report Templates** tab. You can order them by report name, description, subject, date they were last modified, last run, or by whom they were modified.

For more information about the options and actions in the Reports Tab page, see [Report Templates Overview](#).

Table 4-177. Predefined Filter Groups

Filter Group	Description
Name	Filter by the template name. For example, you can list all reports that contain <i>my template</i> in their name by typing my template .
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by those objects.
Owner	Filter by the owner of the report template.

vSphere users must be logged in until the report generation is complete. If you log out or your session expires, the report generation fails.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

Generated Reports Tab

All reports that are generated for a selected object are listed on the **Generated Reports** tab.

In the menu, click **Environment**, and then in the left pane select an object and click **Reports > Generated Reports** to access the Generated Reports tab.

If the report is generated through a schedule, the owner is the user who created the schedule.

Note The maximum number of reports per template is 10. With every new generated report, vRealize Operations Manager deletes the oldest report.

You can filter the reports list by adding a filter from the right side of the panel.

For more information about the options and actions in the Generated Reports tab page, see [Generated Reports Overview](#).

Table 4-178. Predefined Filter Groups

Filter Group	Description
Report Name	Filter by the report template name. For example, you can list all reports that contain <i>my template</i> in their name by typing my template .
Template	Filter by the report template. You can select a template from a list of templates applicable for this object.
Completion Date/Time	Filter by the date, time, or time range.
Status	Filter by the status of the report. On each data node, only one report can be processed. Therefore, reports that are queued can be moved to the processed state only after the previous report on the specific node has failed or completed. The maximum queue time is restricted to 4 hours. After 4 hours, if processing of the report has not started, the report is marked as failed.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by those objects.

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

Create and Modify a Report Template

You create a report to generate a scheduled snapshot of views and dashboards. You can track current resources and predict potential risks to the environment. You can schedule automated reports at regular intervals.

Procedure

- 1 In the menu, click **Dashboards**, and then in the left pane click **Reports**.
- 2 On the **Report Templates** tab, click **Add** to create a template.
- 3 Complete the steps in the left pane to:
 - a Enter a name and description for the report template.
[Name and Description Details](#)
 - b Add a view or a dashboard.
[Views and Dashboards Details](#)
 - c Select an output for the report.
[Formats Details](#)
 - d Select the layout options.
[Layout Options Details](#)
- 4 Click **Save**.

Name and Description Details

The name and description of the report template as they appear in the list of templates on the **Report Templates** tab.

Where You Add Name and Description

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New Template dialog box, in the workspace, on the left, click **Name and Description**.

Table 4-179. Name and Description Options in the Report Template Workspace

Option	Description
Name	Name of the template as it appears on the Report Templates tab.
Description	Description of the template.

Views and Dashboards Details

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs.

Where You Add Views and Dashboards

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New Template dialog box, in the workspace, on the left, click **Views and Dashboards**. If you create a template, complete the required previous steps of the workspace.

How You Add Views and Dashboards

To add a view or a dashboard to your report template, select it from the list on the left pane and drag it to the main panel. You can drag the views and dashboards in the main panel to reorder them. You can select a portrait or landscape orientation for each view or dashboard from the drop-down menu next to its title.

Table 4-180. Views and Dashboards Options in the Report Template Workspace

Option	Description
Data type	Select Views or Dashboards to display a list of available views or dashboards that you can add to the template.
Create View	Create a view directly from the template workspace. This option is available when you select Views from the Data type drop-down menu.
Edit View	Edit a view directly from the template workspace. This option is available when you select Views from the Data type drop-down menu.
Create Dashboard	Create a dashboard directly from the template workspace. This option is available when you select Dashboards from the Data type drop-down menu.
Edit Dashboard	Edit a dashboard directly from the template workspace. This option is available when you select Dashboards from the Data type drop-down menu.
Quick Filter	Search for views or dashboards by name. To see the complete list of views or dashboards, delete the search box contents and press Enter.
List of views	List of the views that you can add to the template. This list is available when you select Views from the Data type drop-down menu.
List of dashboards	List of the dashboards that you can add to the template. This list is available when you select Dashboards from the Data type drop-down menu.
Preview of views and dashboards	In the main panel, you see a preview of the views and dashboards that you add. When you create a template in the context of an object from the environment, you see a live preview of the views and dashboards.
Colorization	You can enable or disable a colorized PDF output for each list view. This option is available from the right panel when you select Views from the Data type drop-down menu.

Formats Details

The formats are the outputs in which you can generate the report.

Where You Add Formats

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New Template dialog box, in the workspace, on the left, click **Formats** to select a format for the report template. If you create a template, complete the required previous steps of the workspace.

Table 4-181. Formats Options in the Report Template Workspace

Option	Description
PDF	With the PDF format, you can read the reports, either on or off line. This format provides a page-by-page view of the reports, as they appear in printed form.
CSV	In the CSV format, the data is in a structured table of lists.

Layout Options Details

The report template can contain layout options such as a cover page, table of contents, and footer.

Where You Add Layout Options

To create or edit report templates, in the menu, click **Dashboards**, and then in the left pane click **Reports**. From the Report Templates tab on the right pane, click **Add** to add a template. From the New Template dialog box, in the workspace, on the left, click **Layout Options**. If you create a template, complete the required previous steps of the template.

Table 4-182. Layout Options in the Report Template Workspace

Option	Description
Cover Page	Can contain an image up to 5 MB. The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page.
Table of contents	Provides a list of the template parts, organized in the order of their appearance in the report.
Footer	Includes the date when the report is created, a note that the report is created by vRealize Operations Manager, and page number.

Add a Network Share Plug-In for vRealize Operations Manager Reports

You add a Network Share plug-in when you want to configure vRealize Operations Manager to send reports to a shared location. The Network Share plug-in supports only SMB version 2.1.

Prerequisites

Verify that you have read, write, and delete permissions to the network share location.

Procedure

- 1 In the menu, click **Administration** and then in the left pane, click **Management > Outbound Settings**.

- 2 From the toolbar, click **Add**.

- 3 From the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.

The dialog box expands to include your plug-in instance settings.

- 4 Enter an **Instance Name**.

This is the name that identifies this instance that you select when you later configure notification rules.

- 5 Configure the Network Share options appropriate for your environment.

Option	Description
Domain	Your shared network domain address.
User Name	The domain user account that is used to connect to the network.
Password	The password for the domain user account.
Network share root	<p>The path to the root folder where you want to save the reports. You can specify subfolders for each report when you configure the schedule publication.</p> <p>You must enter an IP address. For example, <code>\\IP_address\ShareRoot</code>. You can use the host name instead of the IP address if the host name is resolved to an IPv4 when accessed from the vRealize Operations Manager host.</p> <p>Note Verify that the root destination folder exists. If the folder is missing, the Network Share plug-in logs an error after 5 unsuccessful attempts.</p>

- 6 Click **Test** to verify the specified paths, credentials, and permissions.

The test might take up to a minute.

- 7 Click **Save**.

The outbound service for this plug-in starts automatically.

- 8 (Optional) To stop an outbound service, select an instance and click **Disable** on the toolbar.

Results

This instance of the Network Share plug-in is configured and running.

What to do next

Create a report schedule and configure it to send reports to your shared folder. See [Schedule Reports Overview](#).

Report Templates Overview

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs.

In the menu, click **Dashboards**, and then in the left pane select **Reports**. Select the **Report Templates** tab in the right pane.

The listed templates are user-defined and predefined by vRealize Operations Manager. You can order them by template name, description, subject, date they were modified, last run report, or the user who modified them. For each template, you can see the number of generated reports and schedules.

You can filter the reports based on the name of the report template, the subject, and the owner. You can click **Add** to create a report template. For information about creating a report template, see [Create and Modify a Report Template](#).

You can select a report template from the list, click the vertical ellipsis against each report template, and select options such as run, edit, schedule, delete, clone, and export a report.

Table 4-183. Predefined Filter Groups

Filter Group	Description
Name	Filter by the template name. For example, type my template to list all reports that contain the my template phrase in their name.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by the other objects.
Owner	Filter by the owner of the report template.

The maximum number of reports per template is 10. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

Report Template Actions

You can select more than one report template and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

Option	Description
Delete	Deletes the report template.
Export	Downloads the report template.

Option	Description
Import	Allows you to import a report template by selecting a report template in XML or zip file format.
Change default cover image	Allows you to change the default cover image of the report template. For more information, see Upload a Default Cover Page Image for Reports .

Generated Reports Overview

A report is a scheduled snapshot of views and dashboards. It presents data in formats that can be downloaded.

In the menu, click **Dashboards**, and then in the left pane select **Reports**. Select the **Generated Reports** tab in the right pane.

The list contains all generated reports. If the report is generated through a schedule, the owner is the user who created the schedule.

Note The maximum number of reports per template is 10. After the tenth report is generated, vRealize Operations Manager deletes the oldest report.

To select a generated report from the list, click the vertical ellipsis against each generated report and select options such as run and delete. You can also select more than one generated report and select **Delete** from the **Actions** drop-down menu to delete a generated report.

You can filter the reports list by adding a filter from the upper-right corner of the panel.

Table 4-184. Predefined Filter Groups

Filter Group	Description
Report Name	Filter by the report template name. For example, type my template to list all reports that contain the my template phrase in their name.
Template	Filter by the report template. You can select a template from a list of templates applicable for this object.
Completion Date/Time	Filter by the date, time, or time range.
Subject	Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by that second object.
Status	Filter by the status of the report.

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

If you log in to vRealize Operations Manager with vCenter Server credentials and generate a report, the generated report is always blank.

Generate and Regenerate a Report

To generate a report, use a report template.

Prerequisites

Create a report template.

Procedure

- 1 In the menu, click **Environment**.
- 2 In the left pane, navigate to the relevant object.
- 3 Click the **Reports** tab and click **Report Templates**.
The listed report templates are associated with the current object.
- 4 Navigate to the relevant report template, click the vertical ellipsis, and select **Run**.

Results

The report is generated and listed on the **Generated Reports** tab.

Note To regenerate the selected report, from the **Generated Reports** tab, click the vertical ellipsis against the generated report and select **Run**.

What to do next

Download the generated report and verify the output.

Download a Report

To verify that the information appears as expected, you download the generated report.

Prerequisites

Generate a report.

Procedure

- 1 In the menu, click **Environment**.
- 2 In the left pane, navigate to the object for which you want to download a report.
- 3 Click the **Reports** tab and click **Generated Reports**.
The listed reports are generated for the current object.
- 4 Click the PDF or the CSV icon in the Download column to download the report.

Results

vRealize Operations Manager saves the report file.

What to do next

Schedule a report generation and set the email options, so your team receives the report.

Schedule Reports Overview

The schedule of a report is the time and recurrence of a report generation.

Where Do You Schedule a Report

To schedule a report generation, in the menu, click **Environment**, and then in the left pane navigate to an object and click the **Reports** tab. Select a template to schedule, click the vertical ellipsis, and then click **Schedule**. To edit the schedule of a report, click the **Schedules** link of a report from the **Report Templates** tab, and then from the **Scheduled Reports** dialog box, click **Edit Schedule**.

How Do You Schedule a Report

Table 4-185. Schedule Report Options

Option	Description
Recurrence	Schedule a report to run automatically at regular intervals.
Publishing	<p>Email a generated report to a predefined email group or to a network shared location. For more information about how to set up and configure the email options, see Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts.</p> <p>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.</p> <p>You can add a relative path to upload the report to a predefined sub folder of the Network Share Root folder. For example, to upload the report to the share host C:/documents/uploadedReports/SubFolder1, in the Relative Path text box, enter SubFolder1. To upload the report to the Network Share Root folder, leave the Relative Path text box empty.</p>

Note Only users created in vRealize Operations Manager can add and edit report schedules.

Table 4-186. Scheduled Reports Toolbar Options

Options	Description
New Schedule	You can create a schedule for the report.
Edit Schedule	You can edit an existing report schedule.
Delete Schedule	You can delete an existing report schedule.
Transfer Report Schedule	You can assign a new owner for the selected report schedule. You can select a target user from the Transfer Report Schedules dialog box.

Schedule a Report

To generate a report on a selected date, time, and recurrence, you create a schedule for the report template. You set the email options to send the generated report to your team.

The date range for the generated report is based on the time when generates the report and not on the time when you schedule the report or when vRealize Operations Manager places the report in the queue.

Prerequisites

- Download the generated report to verify the output.
- To enable sending email reports, you must have configured Outbound Alert Settings. See [Notifications](#) .

Procedure

- 1 In the menu, click **Environment**.
- 2 In the left pane, navigate to the object.
- 3 Click the **Reports** tab and click **Report Templates**.
- 4 Select the relevant report template from the list.
- 5 Click the vertical ellipsis and select **Schedule**.
- 6 Select the time zone, date, hour, and minutes (in the range of 0, 15, 30, and 45 minutes) to start the report generation.

vRealize Operations Manager generates the scheduled reports in sequential order. Generating a report can take several hours. This process might delay the start time of a report when the previous report takes an extended period of time.

- 7 From the **Recurrence** drop-down menu, select one of the following options for report generation:

Option	Description
Daily	You can set the periodicity in days. For example, you can set report generation to every two days.
Weekly	You can set the periodicity in weeks. For example, you can set report generation to every two weeks on Monday.
Monthly	You can set the periodicity in months.

- 8 Select the **Email report** check box to send an email with the generated report.
 - a In the **Email addresses** text box, enter the email addresses that must receive the report. You can also add email addresses in the CC list and BCC list.
 - b Select an outbound rule.

An email is sent according to this schedule every time a report is generated.

- 9 Save a generated report to an external location.
- 10 You can add a relative path to upload the report to a predefined sub folder of the Network Share Root folder.

To upload the report to the Network Share Root folder, leave the **Relative Path** text box empty.

- 11 Click **OK**.

What to do next

You can edit, clone, and delete report templates. Before you do, familiarize yourself with the consequences of these actions.

When you edit a report template and delete it, all reports generated from the original and the edited templates are deleted. When you clone a report template, the changes that you make to the clone do not affect the source template. When you delete a report template, all generated reports are also deleted.

Upload a Default Cover Page Image for Reports

You can upload a common default image for the cover page of reports. You do not have to upload a cover page for each report. The cover pages of predefined reports are modified when you use this option. The cover pages of user-defined reports do not change.

Where Do You Upload a Default Cover Page Image for Reports

To upload a default cover page for reports, in the menu, click **Environment**, and then in the left pane navigate to an object and click the **Reports** tab. From the **Report Templates** tab, click the horizontal ellipsis next to the **Add** option and click the **Change default cover image** option.

How Do You Upload a Default Cover Page Image for Reports

Browse for the image that you want to add to the cover page and click **Save**. You can also use the default product image that is available.

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 can monitor and when your licenses expire. A new license key is required for vRealize Operations Manager 7.0 and later versions. All license keys except vSOM Enterprise Plus and its add-ons are invalidated. The product works in evaluation mode until a new valid license key is installed. After you log in to the

user interface of vRealize Operations Manager, if you see that you are using an evaluation license, consider applying for a new license before the end of the 60-day evaluation period.

You can obtain the new license keys from the [MyVMware](#) portal.

Note If you added new licenses when you upgraded to vRealize Operations Manager 7.0, you can skip this step. However, if you have deployed a new instance of vRealize Operations Manager 8.x, you must install a new license.

How License Keys Work

License keys activate the solution or product and are available in varying levels. Higher levels typically allow vRealize Operations Manager to monitor more objects.

Where You Find the License Keys

- 1 In the menu, click **Administration**, and in the left pane click **Management > Licensing**.
- 2 Click the **License Keys** tab.

License Key Options

The options include toolbar and data grid options.

Click **Add** or click the **Horizontal Ellipses** to refresh or remove license keys.

Table 4-187. License Key Toolbar Options

Option	Description
Add	Select a solution or product, and then enter and validate a license key for it.
Delete	Remove a license key.
Refresh License Usage	Update the list of keys.

Use the data grid options to view item details.

Table 4-188. License Key Data Grid Options


Option	Description
Product or Solution	Name of the product or solution associated with the key.
License Type	Level of the license. To view the license edition, click the  icon, and then click About . The About vRealize Operations Manager dialog box opens. You can view the version no and the license edition that is in use.
License Capacity	Number of objects that the license allows the product to monitor.

Table 4-188. License Key Data Grid Options (continued)

Option	Description
License Usage	Number of monitored objects that count against the capacity. If you have an unlimited capacity, this number is zero (0).
Status	Indicates whether the license is valid.
Expiry	Date and time when the license expires.
License Information (below)	Details for the selected license key.
Overview	Solution or product, expiration, capacity, type, and use of the selected license key.
Associated License Groups	License groups that this key is a member of, and the number of objects in the groups.

vRealize Operations Manager License Groups

Like other vRealize Operations Manager groups, you create a license group of objects as a way of gathering those objects for data collection. In this case, you are associating the objects with a product license.

How License Groups Work

License groups require that you select one or more keys that you already added for solution or product activation, and add objects as members to a custom group for those licenses. You might, for example, want to add objects into groups that are associated with a particular level of license key, and monitor or manage by level of key in order to control licensing costs.

Where You Find the License Groups

- 1 In the menu, click **Administration**, and then in the left pane click **Management > Licensing**.
- 2 Click the **License Groups** tab.

License Groups

vCloud Suite

Host CPU-based licenses applied to an object type "Host system" for a given set of clusters. When you apply a CPU license to a group containing Hosts, the VMs on the Hosts will still show "License is invalid" watermark.

VM Licenses

VM based licenses applied to an object type "Virtual Machine" for all other VMs except those on hosts licensed with vCloud Suite. When you apply a VM license key to Virtual Machines, the Hosts on which those VMs run will still show the "License is invalid" watermark.

Note In vRealize Operations Manager, it is possible to mix Operating System Instance (OSI) and CPU based licenses. By mixing difference kind of licenses, you will need to perform extra configurations, like creating separate license groups for each type of license keys (one for CPU and one for OSI (VM)). It is recommended that you use non overlapping exclusive Licensing Groups to have the best advantage when you mix OSI (VM) and CPU licensing.

However, in vRealize Operations Manager you cannot mix core and standard license with any other advanced and enterprise licenses.

Dynamic

Use dynamic membership criteria, not static "Always include/exclude" lists to avoid manual maintenance of license groups.

Note When the license is applied to the respective Object type of each License key, the related objects (parent or children) are also going to have to be included in membership for the License Group. License in invalid" watermark appears in vRealize Operations Manager 6.6 and later. For more information, see the following KB article [51556](#).

License Group Options

The license group options include toolbar and data grid options.

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 4-189. 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 4-190. 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 if it has an assigned maintenance schedule

Note When you perform maintenance operations, it is good practice to stop the End Point Operations Management agent and to restart it after the maintenance is complete to avoid unnecessary system overhead.

How Maintenance Schedules Work

Maintenance schedules require that you select the days and time-of-day when updates or other object maintenance occurs. Note that creating a maintenance schedule does not activate the schedule. A maintenance schedule must be part of a policy before the schedule can take effect.

Where You Find the Maintenance Schedules

In the menu, click **Administration**, and then in the left pane click **Configuration > Maintenance Schedules**.

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 4-191. Maintenance Schedule Toolbar Options

Option	Description
Add	Open a window in which you can select the maintenance schedule settings for a new schedule.
Edit	Change the maintenance schedule settings for an existing schedule.
Delete	Remove the selected maintenance schedule.

Manage Maintenance Schedules

Add or edit a maintenance schedule to take an object offline. vRealize Operations Manager does not collect data from an object that is offline.

Where You Find Manage Maintenance Schedules

- 1 In the menu, click **Administration**, and then in the left pane click **Configuration > Maintenance Schedules**.
- 2 Click the plus sign to add a maintenance schedule or the pencil to edit the selected object.

Table 4-192. Manage Maintenance Schedule Add or Edit Options

Option	Description
Schedule Name	Name that describes the maintenance schedule
Time Zone	Time zone in which you are currently located
Days	Number of days the maintenance period covers
Recurrence	Specify a maintenance schedule to run over a selected period <ul style="list-style-type: none"> ■ 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](#). For information about roles and associated permissions, see [KB 59484](#).

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the **Roles** tab.
- 3 Click the **Add** icon on the toolbar to create a role.

The **Create Role** dialog box appears.

- 4 For the role name, type **admin_cluster**, then type a description and click **OK**.

The `admin_cluster` role appears in the list of roles.

- 5 Click the **admin_cluster** role.
- 6 In the Details grid below, on the Permissions pane, click the **Edit** icon.

The **Assign Permissions to Role** dialog box appears.

- 7 Select the **Administrative Access - all permissions** check box.
- 8 Click **Update**.

This action gives this role administrative access to all the features in the environment.

What to do next

Create a user account, and assign this role to the account.

Create a User Account

As an administrator you assign a unique user account to each user so that they can use vRealize Operations Manager. While you set up the user account, you assign the privileges that determine what activities the user can perform in the environment, and upon what objects.

In this procedure, you will create a user account, assign the `admin_cluster` role to the account, and associate the objects that the user can access while assigned this role. You will assign access to objects in the vRealize Operations Cluster. Then, you will test the user account to confirm that the user can access only the specified objects.

Prerequisites

Create a new role. See [Create a New Role](#).

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the **User Accounts** tab.

- 3 Click the **Add** icon to create a new user account, and provide the information for this account.

Option	Description
User Name	Type the user name to use to log in to vRealize Operations Manager.
Password	Type a password for the user.
Confirm Password	Type the password again to confirm it.
First Name	Type the user's first name. For this scenario, type Tom .
Last Name	Type the user's last name. For this scenario, type User .
Email Address	(Optional). Type the user's email address.
Description	(Optional). Type a description for this user.
Disable this user	Do not select this check box, because you want the user to be active for this scenario.
Require password change at next login	Do not select this check box, because you do not need to change the user's password for this scenario.

- 4 Click **Next**.

The list of user groups appears.

- 5 Select a user group to add the user account as a member of the group.
- 6 Click the **Objects** tab.
- 7 Select the **admin_cluster** role from the drop-down menu.
- 8 Select the **Assign this role to the user** check box.
- 9 In the Object Hierarchies list, select the **vRealize Operations Cluster** check box.
- 10 Click **Finish**.

You created a new user account for a user who can access all the vRealize Operations Cluster objects. The new user now appears in the list of user accounts.

- 11 Log out of vRealize Operations Manager.
- 12 Log in to vRealize Operations Manager as Tom User, and verify that this user account can access all the objects in the vRealize Operations Cluster hierarchy, but not other objects in the environment.
- 13 Log out of vRealize Operations Manager.

Results

You used a specific role to assign permission to access all objects in the vRealize Operations Cluster to a user account named Tom User.

What to do next

Import a user account from an external LDAP user database that resides on another machine, and assign permissions to the user account.

Import a User Account and Assign Permissions

You can import user accounts from external sources, such as an LDAP database on another machine, or a single sign-on server, so that you can give permission to those users to access certain features and objects in vRealize Operations Manager.

Prerequisites

- Configure an authorization source. See [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.
Import single sign-on user groups after adding the current source?	Select Yes so that the wizard directs you to the Import User Groups page when you have completed the SSO source setup. If you want to import user accounts, or user groups at a later stage, select No .
Advanced options	If your environment uses a load balancer, enter the IP address of the load balancer.

- 5 Click **Test** to test the source connection, and then click **OK**.

The certificate details are displayed.

- 6 Select the **Accept this Certificate** check box, and click **OK**.
- 7 In the Import User Groups dialog box, import user accounts from an SSO server on another machine.

Option	Action
Import From	Select the single sign-on server you specified when you configured the single sign-on source.
Domain Name	Select the domain name from which you want to import user groups. If Active Directory is configured as the LDAP source in the PSC, you can only import universal groups and domain local groups if the vCenter Server resides in the same domain.
Result Limit	Enter the number of results that are displayed when the search is conducted.
Search Prefix	Enter a prefix to use when searching for user groups.

- 8 In the list of user groups displayed, select at least one user group, and click **Next**.
- 9 In the Roles and Objects pane, select a role from the **Select Role** drop-down menu, and select the **Assign this role to the group** check box.
- 10 Select the objects users of the group can access when holding this role.
To assign permissions so that users can access all the objects in vRealize Operations Manager, select the **Allow access to all objects in the system** check box.
- 11 Click **OK**.
- 12 Familiarize yourself with single-sign on and confirm that you have configured the single sign-on source correctly.
 - a Log out of vRealize Operations Manager.
 - b Log in to the vSphere Web Client as one of the users in the user group you imported from the single sign-on server.
 - c In a new browser tab, enter the IP address of your vRealize Operations Manager environment.
 - d If the single sign-on server is configured correctly, you are logged in to vRealize Operations Manager without having to enter your user credentials.

Edit a Single Sign-On Source

Edit a single sign-on source if you need to change the administrator credentials used to manage the single sign-on source, or if you have changed the host of the source.

When you configure an SSO source, you specify either the IP address or the FQDN of the host machine where the single sign-on server resides. If you want to configure a new host, that is, if the single sign-on server resides on a different host machine than the one configured when the source was set up, vRealize Operations Manager removes the current SSO source, and creates a new source. In this case, you must reimport the users you want to associate with the new SSO source.

If you want to change the way the current host is identified in vRealize Operations Manager, for example, change the IP address to the FQDN and the reverse, or update the IP address of the PSC if the IP address of the configured PSC has changed, vRealize Operations Manager updates the current SSO source, and you are not required to reimport users.

Procedure

- 1 Log in to vRealize Operations Manager as an administrator.
- 2 In the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.
- 3 Select the single sign-on source and click the **Edit** icon.
- 4 Make changes to the single sign-on source, and click **OK**.

If you are configuring a new host, the New Single Sign-On Source Detected dialog box appears.

- 5 Enter the administrator credentials that were used to set up the single sign-on source, and click **OK**.

The current SSO source is removed, and a new one created.

- 6 Click **OK** to accept the certificate.
- 7 Import the users you want to associate with the SSO source.

Access Control in vRealize Operations Manager

Each user must have a unique account with one or more roles assigned to enforce a role-based security when they use vRealize Operations Manager. You create a user account, and assign the account to be a member of one or more user groups to allow the user to inherit the roles and objects associated with the user group.

Where You Find the Access Control Options

You can manage user accounts and their associated user groups, roles, and passwords.

In the menu, click **Administration**, and then click **Access > Access Control**.

Table 4-193. Access Control Tabs

Option	Description
User Accounts	<p>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.</p> <p>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.</p>
User Groups	<p>Add, edit, or remove, or import vRealize Operations Manager user groups, update the members in a group and the associated objects that they can access. Import user groups from an LDAP database or a single sign-on database that resides on another machine.</p> <p>vRealize Operations Manager continuously synchronizes the user membership of imported LDAP user groups when the autosync option is enabled in the LDAP configuration.</p>
Roles	<p>For users to perform actions in vRealize Operations Manager, they must be assigned specific roles. With role-based access, when you assign a role to a user, you are determining not only what actions the user can perform in the system, but also the objects upon which those actions can be performed while holding the role. For example, to import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management.</p>
Password Policy	<p>Manage local user passwords, set the criteria for account lockout, password strength, and the password change policy settings.</p>

Access Control: User Accounts Tab

You can add, edit, or remove vRealize Operations Manager user accounts, and import user accounts from an external LDAP database. With access control, you manage roles, the objects a user can access while assigned a specific role, and the membership in user groups.

Where You Manage User Accounts

In the menu, click **Administration**, and then click **Access > Access Control**.

Table 4-194. Access Control User Accounts Summary Grid

Summary Grid Options	Description
User Accounts toolbar	<p>To manage user accounts, use the toolbar icons.</p> <ul style="list-style-type: none"> ■ Click the Add icon to add a user account, and provide the details for the user account in the Add User Account dialog box. ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Edit the selected user account, and modify the details for the user group in the Edit User Account dialog box. ■ Delete. Delete a user account. ■ Click the Horizontal Ellipses and click Import Users to import a user account from an authentication source.
First Name	User's first name, created when you create the user account.
Last Name	User's last name, created when you create the user account.

Table 4-194. Access Control User Accounts Summary Grid (continued)

Summary Grid Options	Description
User Name	User name, without spaces, that will log in to vRealize Operations Manager
Email	User's email address, created when you create the user account.
Description	Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access privileges.
Imported	Indicates whether the user account is imported or not.
Source Type	Indicates whether the user account is a local user, or an external user who is integrated through an external authentication source, such as from LDAP, SSO, AD, OpenLDAP, vCenter Server.
Enabled	Indicates whether the user account is enabled to use vRealize Operations Manager features. An administrator can edit a user account to manually enable it, or disable it to prevent user access to vRealize Operations Manager.
Locked	Indicates whether vRealize Operations Manager has locked the user account. For example, a user account can get locked based on the password lockout policy, or if the user enters an incorrect password three times in the span of five minutes.
Access All Objects	Indicates whether the user account is allowed to access all the objects that are imported into the vRealize Operations Manager instance.

After you add a user account, use the Details grid to view and edit which user accounts are assigned to user groups, and view the permissions assigned to the user account.

Table 4-195. Access Control User Accounts Details Grid

Details Grid Options	Description
User Groups	<p>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.</p> <ul style="list-style-type: none"> ■ User Name: Identifies the user account. To change the user groups associated with the user account, click the Edit icon. <p>The Choose Groups Membership dialog box opens.</p> <ul style="list-style-type: none"> ■ Click the All tab to view all the available groups. ■ Click the Selected tab to view the groups that the user account is part of. ■ Click the Unselected tab to view the groups that the user account is not a part of. ■ Use the Search field to search for specific groups. ■ Members: Displays the number of users that are assigned to the user group.
Permissions	<p>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.</p> <ul style="list-style-type: none"> ■ Role: Indicates the name of the role or roles assigned to the user. ■ Role Description: Displays the description entered for the role. ■ Object Hierarchy: Displays the name of the object hierarchy assigned to the user while holding this role. ■ Objects: Displays the number of objects included in the hierarchy that the user can access. ■ Association: Indicates if the role and objects are assigned to the selected user, or assigned to a user group to which the user belongs.

Modify User Accounts and Assign Groups and Permissions

You can add user accounts so that users can access the features of vRealize Operations Manager and certain objects in the environment. Or, modify user accounts to change their attributes, disable or lock the accounts, or require them to change their password. After you add user accounts, you can assign them to one or more user groups, and assign roles and objects to the account to specify the actions the user can perform and upon what objects. Assign the administrators role only to specific users who must access objects and perform actions in the entire environment.

Where You Add or Edit User Accounts

- 1 To add a user account, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**
- 2 In the **User Accounts** tab, click **Add**
- 3 To edit a user account, click the vertical ellipsis and select **Edit**.

Table 4-196. Add or Edit Users Accounts- User Details Page

User Details Options	Description
User Name	User name, without spaces to access the vRealize Operations Manager
Password	User's password to access the vRealize Operations Manager instance.

Table 4-196. Add or Edit Users Accounts- User Details Page (continued)

User Details Options	Description
Confirm Password	Confirmation of the user's password.
First Name	User's first name, created when you create the user account.
Last Name	User's last name, created when you create the user account.
Email Address	User's email address, created when you create the user account.
Description	Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access rights.
Disable this user	Disable the user account so that a user cannot access the vRealize Operations Manager instance.
Account is locked out	Indicates that vRealize Operations Manager has locked the user account.
Require password change at next login	Enable users to change their password the next time they log in to the vRealize Operations Manager instance.

- 4 After you enter the user details, click **Next**.

Table 4-197. Add or Edit User Accounts - Assign Groups and Permissions page

Assign Groups Roles, and Objects Options	Description
Groups	Select or deselect the groups associated with the user account. To select or deselect all accounts, click the Group Name check box. You cannot add user accounts to groups that you imported from an LDAP database.
Objects	<p>Roles determine which actions a user can perform in the system. Select a role from the Select Role drop-down menu, and then select the Assign this role to the user check box. You can associate more than one role with the user account.</p> <p>Select which objects the user can access when assigned this role.</p> <ul style="list-style-type: none"> ■ 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. <p>Note The roles and object permissions are interlinked when you assign more than one role to a user. For example, if the user has both, ReadOnly and PowerUser roles, the permissions associated with the PowerUser role will apply, because the PowerUser role includes the permissions associated with the ReadOnly role along with other permissions.</p> <p>If the user has a custom role and the PowerUser role and the permissions of the custom role are not included in the permissions of the PowerUser role, the permissions of both the roles are merged and applied to the user.</p> <p>The same rule (object permissions from different roles are merged) applies to the object hierarchies as well.</p>

Import User Accounts

You can import user accounts so that users can access the features of vRealize Operations Manager and the objects in the environment. After you import user accounts, you can assign them to user groups and roles. You can also specify the objects users can access while using the assigned roles.

Where You Import User Accounts

- 1 To import user accounts, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the horizontal ellipsis next to **Add** and then, click **Import Users**.

Table 4-198. Import Users from a LDAP Source

User Details Options	Description
Import From	LDAP host machine, Active Directory, or Other sources configured to import user accounts. <ul style="list-style-type: none"> ■ Add icon. Add an LDAP import source, and provide the information for the LDAP import source in the Add Source for User and Group Import dialog box. ■ Edit icon. Edit the selected LDAP import source, and modify the details in the Edit Source for User and Group Import dialog box.
User Name	Click Change Credentials to display the user name of the LDAP source credential used to import user accounts to the vRealize Operations Manager instance.
Password	Password for the LDAP source credential to import user accounts to the vRealize Operations Manager instance.
Search String	Enter a search string, and click Search to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the User Name check box to import all users. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

Table 4-199. Import Users from a VMware Identity Manager Source

User Details Options	Description
Import From	VMware Identity Manager configured as the source to import user accounts. <ul style="list-style-type: none"> ■ Add icon. Add a VMware Identity Manager import source, and provide the information for the VMware Identity Manager import source in the Add Source for User and Group Import dialog box. ■ Edit icon. Edit the selected VMware Identity Manager import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Search Prefix	Enter a search string, and click Search to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the User Name check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

Table 4-200. Import Users from a Single Sign On Source

User Details Options	Description
Import From	SSO source configured as the source to import user accounts. <ul style="list-style-type: none"> ■ Add icon. Add an SSO import source, and provide the information for the SSO import source in the Add Source for User and Group Import dialog box. ■ Edit icon. Edit the selected SSO import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Result Limit	Determines the number of users displayed.
Search Prefix	Enter a search prefix, and click Search to start the search for user accounts.
User Name Summary grid	Lists the users available for import. Select the check box for each user to import, or select the User Name check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to vRealize Operations Manager do not appear in the list.

- 3 After you enter the import users details, click **Next**.

Table 4-201. Import Users Accounts- Assign Groups and Permissions Page

Assign Groups Roles, and Objects Options	Description
Groups	Select or deselect the groups associated with the user account. To select or deselect all accounts, click the Group Name check box. You cannot add user accounts to groups imported from LDAP.
Objects	Select or deselect roles in the Select Role drop-down menu. When you have selected a role, click the Assign this role to the user check box. You can assign more than one role to a user account. Select which objects the user can access when assigned this role. <ul style="list-style-type: none"> ■ Select Object Hierarchies: Displays groups of objects. Select an object in this list to select all the objects in the hierarchy, ■ Select Object: To select specific objects within the object hierarchy, click the down arrow to expand the list of objects. For example, expand the Adapter Instance hierarchy, and select one or more adapters. ■ Allow access to all objects in the system: Select this check box to permit the user account access to all objects in the system.

Access Control: User Groups Tab

You can manage the user groups associated with the users and objects in your environment. You can import user groups from an LDAP database that resides on another machine, or from a single sign-on server.

Where You Manage User Groups

- 1 To manage user groups, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Click the **User Groups** tab.

Table 4-202. Access Control User Groups Summary Grid

Option	Description
User Groups toolbar	<p>To manage user groups, use the toolbar icons.</p> <ul style="list-style-type: none"> ■ Click the Add icon to add a user group, and provide the details for the user group in the Add User Group dialog box. ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Edit the selected user group, and modify the details for the user group in the Edit User Group dialog box. ■ Clone. Clone a user group, and enter a name and description for the cloned user group. ■ Delete. Delete a user group. ■ Click the Horizontal Ellipses and click Import to import a user group, and provide the details to import the user group in the Import User Groups dialog box.
Group Name	Name of the user group.
Description	Description of the group, indicating its purpose.
Members	Number of members in the group.
Group Type	Type of group, either a local user group or a group imported from LDAP.
Distinguished Name	Names for LDAP objects, such as domains and users.
Access All Objects	Indicates if the user group account is allowed to access all the objects that are imported into the vRealize Operations Manager instance.

After you select a user group in the summary grid, view details about associated users in the Details pane.

Table 4-203. Access Control User Groups Details Grid

Option	Description
User Accounts	<p>Associated user accounts appear when you click a user group in the summary grid. You can then view or modify user accounts that are part of the selected group.</p> <ul style="list-style-type: none"> ■ User Name: Name of each user who is a member of the selected group. To change the user accounts associated with the user group, click the Add icon. <p>The Add Users to Group dialog box opens.</p> <ul style="list-style-type: none"> ■ Click the All tab to view all the available user accounts. ■ Click the Selected tab to view the user accounts that are part of the group. ■ Click the Unselected tab to view the user accounts that are not a part of the group. ■ Use the Search field to search for specific user accounts. ■ First Name: First name of each user account in the group. ■ Last Name: Last name of each user account in the group. <p>You can remove a user from the group by selecting the user in the Details pane and clicking Delete</p>
Permissions	<p>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.</p> <ul style="list-style-type: none"> ■ Role Name: Indicates the roles assigned to the selected user group. ■ Role Description: Description for the selected user group, defined when you created the group. ■ Object Hierarchy: The names of the object hierarchies assigned to the group while holding a specific role. ■ Objects: The number of objects the user group can access within the selected hierarchy.

Add User Groups and Assign Members and Permissions

You can view and modify the details for user groups, including users, roles, and objects.

Where You Add User Groups

- 1 To add a user group, in the menu, click **Administration** and then click **Access > Access Control**.
- 2 Select the **User Groups** tab and then click the **Add** icon.

Table 4-204. Add or Edit User Group - Name and Description Page

Option	Description
Group Name	Name of the user group, either created manually, imported from a single sign-on server, or imported from an LDAP database that resides on another machine.
Description	Description of the user group, indicating its purpose.

- 3 After you enter the name and description, click **Next**

Table 4-205. Add or Edit User Group - Assign Members and Permissions Page

Option	Description
Members	Select the members associated with the user group.
Objects	<p>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.</p> <p>Select which objects the users of the group can access when assigned this role.</p> <ul style="list-style-type: none"> ■ 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. <p>Note The roles and object permissions are interlinked when you assign more than one role to a user. For example, if the user has both, ReadOnly and PowerUser roles, the permissions associated with the PowerUser role will apply. The PowerUser role includes the permissions associated with the ReadOnly role along with other permissions.</p> <p>If the user has a custom role and the PowerUser role and the permissions of the custom role are not included in the permissions of the PowerUser role. The permissions of both the roles are merged and applied to the user.</p> <p>The same rule (object permissions from different roles are merged) applies to the object hierarchies as well.</p>

Import User Groups

You import user groups from a single sign-on server, VMware Identity Manager, Active Directory, or an LDAP database on another machine so that you can use those groups in vRealize Operations Manager.

Where You Import User Groups

- 1 To import a user group, in the menu, click **Administration**, and then in the left pane click **Access > Access Control**.
- 2 Select the **User Groups** tab and click the **Import Group** icon.

The options displayed in the Import User Groups page depend upon the authentication source you select.

Table 4-206. Import User Groups Page - LDAP, Active Directory, and Others Sources

Option	Description
Import From	Host machine configured as the source to import the user groups. These options are displayed when the host machine of an LDAP, Active Directory, or Other source is selected.
User Name	User name of the source credential to import user groups to the vRealize Operations Manager instance.

Table 4-206. Import User Groups Page - LDAP, Active Directory, and Others Sources (continued)

Option	Description
Password	Password for the source credential to import user groups to the vRealize Operations Manager instance.
Search String	Invoke the search for user groups.
Advanced	<p>Displays the advanced import settings.</p> <ul style="list-style-type: none"> ■ Group Search Criteria. Search criteria to find LDAP groups. If not included, vRealize Operations Manager uses the default search parameters: <code>((objectClass=group) (objectClass=groupOfNames))</code> ■ Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, vRealize Operations Manager uses member by default. ■ User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form <code>((key1=value1) (key2=value2))</code>. If not included, vRealize Operations Manager searches for each user separately. This operation might take extra time. ■ Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, vRealize Operations Manager treats the member entry as a distinguished name. ■ LDAP Context Attributes. Attributes that vRealize Operations Manager applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as <code>java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse.</code>
Group Name	Displays the user groups found. Click the check box for each user group to import.

Table 4-207. Import User Groups Page - Single Sign On Source

Option	Description
Import From	Host machine configured as the source to import the user groups.
Domain Name	User name of the source credential to import user groups to the vRealize Operations Manager instance.
Result Limit	Determines the number of groups displayed.
Search Prefix	Enter a search prefix to narrow your search.
Group Name	Displays a list of user groups. Select the Group Name check box to import all the displayed user groups, or select the check box next to each user group that you want to import.

Table 4-208. Import User Groups from a VMware Identity Manager Source

User Details Options	Description
Import From	<p>VMware Identity Manager configured as the source to import user groups.</p> <ul style="list-style-type: none"> ■ Add icon. Add an VMware Identity Manager import source, and provide the information for the VMware Identity Manager import source in the Add Source for User and Group Import dialog box. ■ Edit icon. Edit the selected VMware Identity Manager import source, and modify the details in the Edit Source for User and Group Import dialog box.
Domain Name	Enter the domain name for import.
Search Prefix	Enter a search string, and click Search to start the search for user groups.
User Name Summary grid	Lists the users available for import. Select the check box for each user group to import, or select the Group Name check box to import all groups. User groups that are already imported to vRealize Operations Manager do not appear in the list.

- 3 After you enter the import user group details, click **Next**.

Table 4-209. Import User Groups - Roles and Objects Page

Option	Description
Select Role	Displays available roles in a drop-down menu.
Assign this role to the group	Roles determine which actions users of the group can perform in the system. Select a role from the Select Role drop-down menu, and then select the Assign this role to the user check box. You can associate more than one role with the user group.
Select Object Hierarchies	<p>Select which objects the users of the group can access when assigned this role.</p> <ul style="list-style-type: none"> ■ 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 4-210. Access Control Roles Summary Grid

Option	Description
Roles toolbar	<p>To manage roles, use the toolbar icons.</p> <ul style="list-style-type: none"> ■ Click the Add icon. to add a user role, and provide the name and description for the role in the Create Role dialog box. ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Edit the selected user role, and modify the details for the role in the Edit Role dialog box. ■ Clone. Clone the selected user role ■ Delete. Delete a user role.
Role Name	Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions.
Role Description	Description of the role, indicating its purpose.

You can view details for the user accounts and user groups associated with a selected role in the Details panes.

Table 4-211. Access Control Roles Details Panes

Option	Description
User Accounts	<p>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.</p> <ul style="list-style-type: none"> ■ First Name. Indicates the first name of each user who is assigned this role. ■ Last Name. Indicates the last name of each user who is assigned this role. ■ User name , without spaces, that will log in to vRealize Operations Manager ■ Email. Indicates the email address for each user who is assigned this role.
User Groups	<p>The user groups assigned the selected role.</p> <ul style="list-style-type: none"> ■ Group Name: Name of each group that is associated with the selected role. ■ Members: Number of members in each group.
Permissions	<p>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.</p> <p>You can edit the permissions assigned to the role by clicking the Edit icon.</p> <ul style="list-style-type: none"> ■ Click the Expand All button to expand the trees of all three categories, and select the check boxes to apply permissions for the selected role. ■ To assign all the available permissions to the selected role, select the Administrative Access - all permissions check box.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`

- Set CPU Count for VM Power Off Allowed
- Set CPU Count and Memory for VM Power Off Allowed

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

Access Control: Password Policy Tab

To ensure security in vRealize Operations Manager, you must manage user passwords. Determine the criteria used for account lockout, password strength, and the password change policy. When a user session becomes inactive for 30 minutes, the session times out, and the user must log in to vRealize Operations Manager again.

Where You Manage the Password Policy

- 1 To manage user roles, in the menu, click **Administration**, and then click **Access > Access Control**.
- 2 Click **Password Policy** tab.

Account Lockout

Indicates whether the account lockout is in effect, and indicates the number of login attempts allowed before the account is locked. The account lockout policy is enabled by default.

Password Strength

Indicates whether the policy that requires users to strengthen their password is in effect, and the minimum number of characters required to make a strong password. The password strength policy is enabled by default.

Password Change

Indicates whether the policy that requires users to change their password is in effect, how often the password expires, and whether users will receive a warning. The account password change policy is enabled by default.

Modify the Password Policy

You can modify the password policy by clicking **Edit**.

Table 4-212. Access Control Edit Password Policy Settings

Option	Description
Account Lockout	<p>Modify the settings to lock user accounts.</p> <ul style="list-style-type: none"> ■ 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	<p>Modify the settings required for users to create strong passwords.</p> <ul style="list-style-type: none"> ■ Activate Password Strength Policy. When selected, enables the policy to require users to strengthen their password. ■ Minimum password length. Indicates the number of characters required for user passwords. The default length is eight characters. ■ Passwords must contain numbers. Users must include a combination of letters and numbers. ■ Passwords must not match user names. To ensure security, users are not allowed to use their user name as their password. ■ Passwords must contain at least one uppercase and one lowercase letter. When selected, users must include one or more uppercase characters. ■ Passwords must contain special characters. When selected, users must include one or more special characters. Special characters include: !@#\$\$%^&*+=
Password Change	<p>Modify the settings required for users to change their password.</p> <ul style="list-style-type: none"> ■ 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.

Access Control: Login Message Tab

To provide support for Security Technical Implementation Guide (STIG), you can add a Standard Mandatory DoD Notice and Consent Banner for the users who access vRealize Operations Manager. Use the login message tab to set a message that requires an explicit consent before logging in to vRealize Operations Manager.

- 1 To set a login message, in the menu, click **Administration**, and then click **Access > Access Control**.
- 2 Click the **Login Message** tab.
- 3 To enable the login message, click **Edit** and click the **Display on Login** checkbox.

- 4 Enter the **Title** and enter the content you want to display.

Note You can add text and images copied from an external source and edit it using the formatting options available.

- 5 Enter the button label for users to click to provide their consent. The label **Agree** is entered by default.
- 6 Use the **Live Preview** section to view how the message will appear on the Login screen.
- 7 Click **Save**.

Authentication Sources

vRealize Operations Manager uses authentication sources that enable you to import and authenticate users and user group information that reside on another machine: the Lightweight Directory Access Protocol (LDAP) platform-independent protocol, Active Directory, VMware Identity Manager, Single Sign-On, and Others.

Where You Manage Authentication Sources

To manage authentication sources, in the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.

Table 4-213. Authentication Sources Toolbar and Data Grid

Option	Description
Authentication Sources toolbar	<p>To manage authentication sources, use the toolbar icons.</p> <ul style="list-style-type: none"> ■ Add icon: Add an authentication source, and provide the information for the source in the Add Source for User and Group Import dialog box. ■ Click the Vertical Ellipses to perform any one of the following actions: <ul style="list-style-type: none"> ■ Edit. Edit the selected authentication source, and modify the details in the Edit Source dialog box. ■ Delete. Delete an authentication source. ■ Synchronize User Groups. Synchronize users within the groups imported through the selected Active Directory or LDAP authentication source.
Source Display Name	Name that you assign to the authentication source.
Source Type	<p>Indicates the type of directory services access technology to access the source machine where the authentication database of user accounts resides. Options include:</p> <ul style="list-style-type: none"> ■ Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts. ■ Active Directory or Other: Specifies any other LDAP-based directory services, such as Novel or Open DJ, used to import user accounts from an LDAP database on a Linux Mac machine. ■ SSO SAML: An open-standard data format that enables Web browser single sign-on. ■ VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources.
Host	Name or IP address of the host machine where the user database resides.

Table 4-213. Authentication Sources Toolbar and Data Grid (continued)

Option	Description
Port	Port used for the import.
Base DN	Base distinguished name for the user search. vRealize Operations Manager locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.
Auto Synchronization	When selected, enables vRealize Operations Manager to map imported LDAP users to user groups.
Last Synchronized	Date and time that the synchronization last occurred.

Authentication Sources: Add Authentication Source for User and Group Import

When you import user account information that resides on another machine, you must define the criteria used to import the user accounts from the source machine.

Where You Add or Edit Authentication Sources

- 1 To add authentication sources, in the menu, click **Administration**, and then in the left pane click **Access > Authentication Sources**.
- 2 Click **Add**.
- 3 To edit authentication sources, click **Edit**.

Table 4-214. Authentication Sources Add Source for User and Group Import

Option	Description
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 database of user accounts resides. There are two types of databases: LDAP and single sign-on. Options include: <ul style="list-style-type: none"> ■ SSO SAML: An XML-based standard for a web browser single sign-on that enables users to perform single sign-on to multiple applications. ■ Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts. ■ Other: Specifies any other LDAP-based directory services, such as Novel or OpenDJ, used to import user accounts from an LDAP database on a Linux Mac machine. ■ VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources.

Table 4-215. 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.
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. <ul style="list-style-type: none"> ■ 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.

Table 4-215. Authentication Sources Add Source for User and Group Import - Options Available When SSO SAML Is Selected. (continued)

Name	Description
Automatically redirect to vRealize Operations single sign-on URL?	<p>After you have configured a single sign-on source, users are redirected to the vCenter SSO server.</p> <ul style="list-style-type: none"> ■ 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?	<p>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.</p> <ul style="list-style-type: none"> ■ If you select Yes, the wizard directs you to the Import User Groups page so that you can import user groups when you have finished setting up the SSO source. ■ If you want to import user accounts, or user groups at a later stage, select No.
Advanced	If your system uses a load balancer, enter the IP address of the load balancer.
Test	Tests whether the host machine can be reached with the credentials provided.

Table 4-216. 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	<p>Applies basic settings to integrate the LDAP import source with the instance of vRealize Operations Manager.</p> <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> ■ Domain/Subdomain. Domain information for the LDAP user account. ■ Use SSL/TLS. When selected, vRealize Operations Manager uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, vRealize Operations Manager prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds. ■ If Active Directory uses a self-signed certificate, then the certificate should contain the Subject Alternative Name field. vRealize Operations Manager can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the Subject Alternative Name field matches the address of the domain controller on which the certificate is used. ■ User Name. Name of the user account that can log in to the LDAP host machine. ■ Reset Password. Reset the password of the user account that can log in to the LDAP host machine. ■ Automatically synchronize user membership for configured groups. When selected, enables vRealize Operations Manager to map imported LDAP users to user groups. ■ Host. Name or IP address of the host machine where the LDAP user database resides. ■ Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS. ■ Base DN. Base distinguished name for the user search. vRealize Operations Manager locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration. ■ Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is <i>userPrincipalName</i>.
Integration Mode Advanced settings	<p>Applies advanced settings to integrate the LDAP import source with the instance of vRealize Operations Manager.</p>

Table 4-216. Authentication Sources Add Source for User and Group Import - Options Available When Open LDAP, Active Directory, and Other Are Selected. (continued)

Option	Description
	<p>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.</p> <ul style="list-style-type: none"> ■ Host. Name or IP address of the host machine where the LDAP user database resides. ■ Use SSL/TLS. When selected, vRealize Operations Manager uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, vRealize Operations Manager prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds. ■ If Active Directory uses a self-signed certificate, then the certificate should contain the Subject Alternative Name field. vRealize Operations Manager can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the Subject Alternative Name field matches the address of the domain controller on which the certificate is used. ■ Base DN. Base distinguished name for the user search. vRealize Operations Manager will locate only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although vRealize Operations Manager populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration. ■ User Name. Name of the user account that can log in to the LDAP host machine. ■ Reset Password. Reset the password of the user account that can log in to the LDAP host machine. ■ Automatically synchronize user membership for configured groups. When selected, enables vRealize Operations Manager to map imported LDAP users to user groups. ■ Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is <i>userPrincipalName</i>. ■ Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS.

Table 4-216. Authentication Sources Add Source for User and Group Import - Options Available When Open LDAP, Active Directory, and Other Are Selected. (continued)

Option	Description
Search Criteria	<p>Displays the search criteria settings.</p> <p>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.</p> <ul style="list-style-type: none"> ■ Group Search Criteria. Search criteria to find LDAP groups. If not included, vRealize Operations Manager uses the default search parameters: <code>((objectClass=group)(objectClass=groupOfNames))</code> ■ Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, vRealize Operations Manager uses member by default. ■ User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form <code>((key1=value1)(key2=value2))</code>. If not included, vRealize Operations Manager searches for each user separately. This operation might take extra time. ■ Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, vRealize Operations Manager treats the member entry as a distinguished name. ■ LDAP Context Attributes. Attributes that vRealize Operations Manager applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as <code>java.naming.referral=ignore,java.naming.ldap.deleteRDN=false</code>.
Test	<p>Tests whether the host machine can be reached, with the credentials provided.</p> <p>Although a test of the connection is successful, users who use the search feature must have read permissions in the LDAP source.</p> <p>This test does not verify the accuracy of the Base DN or Common Name entries.</p>

Table 4-217. Authentication Sources Add Source for User and Group Import - Options Available When VMware Identity Manager Is Selected.

Option	Description
Host	Name or IP address of the VMware Identity Manager machine where the single sign-on user server resides.
Port	The single sign-on listening port. By default this is set to 443.
Tenant	This is an optional field.
User name	VMware Identity Manager system-domain tenant administrator user name.
Password	Password of the VMware Identity Manager system-domain tenant administrator.

Table 4-217. Authentication Sources Add Source for User and Group Import - Options Available When VMware Identity Manager Is Selected. (continued)

Option	Description
Redirect IP/ FQDN	<p>This is the IP address of vRealize Operations Manager node where a user is redirected after a successful authentication from VMware Identity Manager. By default, this is the IP address of the vRealize Operations Manager primary node.</p> <hr/> <p>Note When the primary replica becomes the primary node on vRealize Operations Manager, then vRealize Operations Manager administrator has to manually edit the IP address and set it to the IP address of the current primary node.</p> <hr/>
Test	Tests whether the VMware Identity Manager machine can be reached, with the credentials provided.

Audit Users and the Environment in vRealize Operations Manager

At times, you might need to provide documentation as an evidence of the sequence of activities that took place in your vRealize Operations Manager environment. Auditing allows you to view the users, objects, and information that is collected. To meet audit requirements, such as for business critical applications that contain sensitive data that must be protected, you can generate reports on the activities of your users, the privileges assigned to users to access objects, and the counts of objects and applications in your environment.

Auditing reports provide traceability of the objects and users in your environment.

User Activity Audit

Run this report to understand the scope of user activities, such as logging in, actions on clusters and nodes, changes to system passwords, activating certificates, and logging out.

User Permissions Audit

Generate this report to understand the scope of user accounts and their roles, access groups, and access privileges.

System Audit

Run this report to understand the scale of your environment. This report displays the counts of configured and collecting objects, the types and counts of adapters, configured and collecting metrics, super metrics, applications, and existing virtual environment objects. This report can help you determine whether the number of objects in your environment exceeds a supported limit.

System Component Audit

Run this report to display a version list of all the components in your environment.

Reasons for Auditing Your Environment

Auditing in vRealize Operations Manager helps data center administrators in the following types of situations.

- You must track each configuration change to an authenticated user who initiated the change or scheduled the job that performed the change. For example, after an adapter changes an object, which is associated with a specific object identifier at a specific time, the data center administrator can determine the principal identifier of the authenticated user who initiated the change.
- You must track who made changes to your data center during a specific range of time, to determine who changed what on a particular day. You can identify the principal identifiers of authenticated users who were logged in to vRealize Operations Manager and running jobs, and determine who initiated the change.
- You must determine which objects were affected by a particular user during a time-specific range of time.
- You must correlate events that occurred in your data center, and view these events overlaid so that you can visualize relationships and the cause of the events. Events can include login attempts, system start up and shutdown, application failures, watchdog restarts, configuration changes of applications, changes to security policy, requests, responses, and status of success.
- You must validate that the components installed in your environment are running the latest version.

User Activity Audit

The user activity report helps you understand the scope of user activities in your vRealize Operations Manager instance, such as when users logged in, actions they took on clusters and nodes, changes they made to system passwords, when they activated certificates, and when they logged out.

Where You Audit User Activity

To audit user activity, in the menu, click **Administration**, and then in the left pane click **History > Audit**. The activities that users performed in the environment appear on the page.

Table 4-218. User Activity Audit Actions

Option	Description
Download	Download the user activity audit information to a report in PDF or XLS format.
Configure	<p>Configure the settings to send the user activity log to an external syslog server to meet security auditing requirements.</p> <ul style="list-style-type: none"> ■ Output log to external syslog server. When selected, vRealize Operations Manager sends the log to a separate server machine. ■ IP Address or Host Name. Identification for the syslog server. ■ Port. vRealize Operations Manager port used to send the audit information to the external server.

Table 4-218. User Activity Audit Actions (continued)

Option	Description
Date Range	Display the list of user activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times.
Starting Line	Indicates the starting line of the file. 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end.
Number of Lines	Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0.
Filter	Filters the data according to User ID, User Name, Auth Source, Session, Message, and Category.

User Permissions Audit

A user permissions audit report provides an overview of the local users and LDAP imported users in your vRealize Operations Manager instance, and a list of groups to which each user belongs. This report helps you understand the scope of the user accounts and their roles, access groups, and access privileges in your environment.

The report displays the access group associated with each local user and LDAP imported user and the access privileges granted to the user in each access group. This report does not include vCenter Server users, roles, or privileges.

The report displays the access group associated with each local user and the access privileges granted to the user in each access group. This report does not include vCenter Server users, roles, or privileges.

When a user is a member of a specific user group, the associated access group could provide the user with access to configuration, dashboards, and templates, or to specific navigation areas in the user interface such as Administration. The access rights associated with the access group include actions for each access group, such as the ability to add, edit, or delete dashboards, or to view, configure, or manage objects.

Where You Audit User Permissions

- 1 To audit user permissions, in the menu, click **Administration**, and then in the left pane click **History > Audit**.
- 2 Click the **User Permissions Audit** tab.

The permissions assigned to users, and their associated access groups and access privileges, appear on the page.

Table 4-219. 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 4-220. 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 4-221. 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 4-222. User Preference Settings

Option	Description
Display	<p>Configure how many metrics and root cause groups to display.</p> <ul style="list-style-type: none"> ■ Color scheme: Set the user interface to display in light or dark colors. ■ Important metrics count to show. Set the number of metrics to display. ■ Root cause groups count to show. Set the number of root cause groups to display. ■ Font. Select the font for reports.
Time	<p>Synchronize the time used for the vRealize Operations Manager instance, and display the updated time when vRealize Operations Manager communicates with the host machine.</p> <ul style="list-style-type: none"> ■ Browser time. All dates and times displayed in the user interface use the time zone settings of the local browser. ■ Host time. All dates and times displayed in the user interface use the time zone of the host machine. ■ Show update time in the application header. Displays the updated time in the top-level header of the vRealize Operations Manager user interface. The updated timestamp appears to the left of the refresh button. Other features, such as dashboards, use the updated time to display data at specific intervals.
Account	Change the password for the user account.

vRealize Operations Manager Passwords and Certificates

For secure vRealize Operations Manager operation, you might need to perform maintenance on passwords or authentication certificates.

- Passwords are for user access to the product interfaces or to console sessions on cluster nodes.
- Authentication certificates are for a secure machine-to-machine communication within vRealize Operations Manager itself or between vRealize Operations Manager and other systems.

Reset the vRealize Operations Manager Administrator Password

You might need to reset the vRealize Operations Manager administrator password as part of securing or maintaining your deployment and if you forget the admin account password.

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at `https://<master-node-name> or <master-node-ip-address>/admin`.
- 2 Log in with the admin user name and password for the master node.
- 3 In the left pane, click **Administrator Settings**.

- 4 In the **Change Administrator Password** section, enter the current password, and enter the new password twice to ensure its accuracy.

Note You cannot change the administrator user name.

- 5 Click **Save**.
- 6 Optionally, to recover a forgotten password, configure the **Password Recovery Settings**.

Table 4-223. Password Recovery Settings

Password Recovery Settings Options	Description
Your E-mail	Email id to which you want to receive the recovery email.
SMTP Server	DNS name or IP address of the SMTP server that is used to send the password recovery email.
Port	Port used for the communication. By default, 25 is used for a non-secure port and 465 for a secure port.
SSL (SMTPS)	Enable or disable to protect the communication using the secure socket layer.
STARTTLS Encryption	Enable or disable to switch the insecure communication starting with the TLS handshake.
Sender E-mail	The email id from which the password recovery email is sent.
User name	User name for the SMTP server account, as some servers require authentication.
Password	Password for the SMTP server account.
Test	To verify the mandatory fields and make an attempt to communicate with the given SMTP server.

- 7 Click **Save**. Optionally, click **Reset** to enter the details again.

Reset the vRealize Operations Manager Administrator Password on vApp Clusters

You must reset the password if the admin account password is lost.

When the vRealize Operations Manager password for the built-in admin account is lost, follow these steps to reset it on vApp clusters.

Prerequisites

This procedure requires root account credentials.

- In vRealize Operations Manager vApp deployments, when you log in to the console of the virtual application for the first time, you are forced to set a root password.

- The vRealize Operations Manager console root password can be different than the admin account password that you set when configuring the vRealize Operations Manager primary node.

Procedure

- 1 Log in to the master node command-line console as `root`.
- 2 Enter the following command, and follow the prompts.

```
$VMWARE_PYTHON_BIN $VCOPS_BASE/../../vmware-vcopssuite/utilities/
sliceConfiguration/bin/vcopsSetAdminPassword.py --reset
```

Reset the vRealize Operations Manager Administrator Password on Windows Clusters

If the admin account password is lost, you must reset the password.

When the vRealize Operations Manager password for the built-in admin account is lost, follow these steps to reset it on Windows clusters.

Procedure

- 1 Open the command prompt using the **Run as Administrator** option.
- 2 Enter the following command, and follow the prompts.

```
%VMWARE_PYTHON_BIN% %VCOPS_BASE%\..\vmware-
vcopssuite\utilities\sliceConfiguration\bin\vcopsSetAdminPassword.py --reset
```

Generate a vRealize Operations Manager Passphrase

When users need to add a node to the vRealize Operations Manager cluster, you can generate a temporary passphrase instead of giving them the primary administrator login credentials, which might be a security issue.

A temporary passphrase is good for one use only.

Prerequisites

Create and configure the primary node.

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at <https://master-node-name-or-ip-address/admin>.
- 2 Log in with the admin user name and password for the master node.
- 3 In the list of cluster nodes, select the master node.
- 4 From the toolbar above the list, click the option to generate a passphrase.
- 5 Enter a number of hours before the passphrase expires.

6 Click **Generate**.

A random alphanumeric string appears, which you can send to a user who needs to add a node.

What to do next

Have the user supply the passphrase when adding a node.

Custom vRealize Operations Manager Certificates

By default, vRealize Operations Manager includes its own authentication certificates. The default certificates cause the browser to display a warning when you connect to the vRealize Operations Manager user interface.

Your site security policies might require that you use another certificate, or you might want to avoid the warnings caused by the default certificates. In either case, vRealize Operations Manager supports the use of your own custom certificate. You can upload your custom certificate during the initial primary node configuration or later.

Custom vRealize Operations Manager Certificate Requirements

A certificate used with vRealize Operations Manager must conform to certain requirements. Using a custom certificate is optional and does not affect vRealize Operations Manager features. You can also use wildcard certificates in vRealize Operations Manager.

Requirements for Custom Certificates

Custom vRealize Operations Manager certificates must meet the following requirements.

- The certificate file must include the terminal (leaf) server certificate, a private key, and all issuing certificates if the certificate is signed by a chain of other certificates.
- In the file, the leaf certificate must be first in the order of certificates. After the leaf certificate, the order does not matter.
- In the file, all certificates and the private key must be in PEM format. vRealize Operations Manager does not support certificates in PFX, PKCS12, PKCS7, or other formats.
- In the file, all certificates and the private key must be PEM-encoded. vRealize Operations Manager does not support DER-encoded certificates or private keys.

PEM-encoding is base-64 ASCII and contains legible BEGIN and END markers, while DER is a binary format. Also, file extension might not match encoding. For example, a generic `.cer` extension might be used with PEM or DER. To verify encoding format, examine a certificate file using a text editor.

- The file extension must be `.pem`.
- The private key must be generated by the RSA or DSA algorithm.
- The private key can be encrypted by a pass phrase. The generated certificate can be uploaded using the primary node configuration wizard or the administration interface.

- The REST API in this vRealize Operations Manager release supports private keys that are encrypted by a pass phrase. Contact VMware Technical Support for details.
- The vRealize Operations Manager Web server on all nodes have the same certificate file, so it must be valid for all nodes. One way to make the certificate valid for multiple addresses is with multiple Subject Alternative Name (SAN) entries.
- SHA1 certificates create browser compatibility issues. Therefore, ensure that all certificates that are created and being uploaded to vRealize Operations Manager are signed using SHA2 or newer.
- The vRealize Operations Manager supports custom security certificates with key length up to 8192 bits. An error is displayed when you try to upload a security certificate generated with a stronger key length beyond 8192 bits.

For more information, see the following KB articles:

- [vRealize Operations Manager 6.x fails to accept and apply Custom CA Certificate \(2144949\)](#)

Configure a Custom Certificate

You can use OpenSSL to configure an authentication certificate for use with vRealize Operations Manager. You must first generate a Certificate PEM for vRealize Operations Manager, then install the Certificate PEM in vRealize Operations Manager. The certificates applied through the vRealize Operations Manager Admin UI will be used only for securely connecting and serving the user interfaces to (external) clients. We do not update the SSL certificates used for establishing a secure connection from vRealize Operations Manager to other services like VMware Identity Manager, vCenter Server, and vRealize Log Insight.

Procedure

- 1 Generate a Certificate PEM file for use with vRealize Operations Manager.

- a Generate a key pair by running this command:

```
openssl genrsa -out key_filename.key 2048
```

- b Use the key to generate a certificate signing request by running this command:

```
openssl req -new -key key_filename.key -out certificate_request.csr
```

- c Submit the CSR file to your Certificate Authority (CA) to obtain a signed certificate.
 - d From your Certificate Authority, download the certificate and the complete issuing chain (one or more certificates). Download them in Base64 format.

- e Enter the command to create a single PEM file containing all certificates and the private key. In this step, the example certificate is *server_cert.cer* and the issuing chain is *cacerts.cer*.

Note The order of CA's certs in the .PEM file: Cert, Private Key, Intermediate Cert and then Root Cert.

`cat server_cert.cer key_filename.key cacerts.cer > multi_part.pem`

In Windows replace `cat` with `type`.

The finished PEM file should look similar to the following example, where the number of CERTIFICATE sections depends on the length of the issuing chain:

```
-----BEGIN CERTIFICATE-----
(Your Primary SSL certificate: your_domain_name.crt)
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
(Your Private Key: your_domain_name.key)
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
(Your Intermediate certificate: DigiCertCA.crt)
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
(Your Root certificate: TrustedRoot.crt)
-----END CERTIFICATE-----
```

2 Install a PEM in vRealize Operations Manager.

- a In a Web browser, navigate to the vRealize Operations Manager administration interface.

```
https://vroops-node-FQDN-or-ip-address/admin
```

- b Log in with the admin user name and password.
- c At the upper right, click the yellow **SSL Certificate** icon.
- d In the **SSL Certificate** window, click **Install New Certificate**.
- e Click **Browse** for certificate.
- f Locate the certificate .pem file, and click Open to load the file in the **Certificate Information** text box. The certificate file must contain a valid private key and a valid certificate chain.
- g Click **Install**.

Verifying a Custom vRealize Operations Manager Certificate

When you upload a custom certificate file, the vRealize Operations Manager interface displays summary information for all certificates in the file.

For a valid custom certificate file, you should be able to match issuer to subject, issuer to subject, back to a self-signed certificate where the issuer and subject are the same.

In the following example, OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32 is issued by OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32, which is issued by OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84, which is issued by itself.

```
Thumbprint: 80:C4:84:B9:11:5B:9F:70:9F:54:99:9E:71:46:69:D3:67:31:2B:9C
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32
Subject Alternate Name:
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:24.000Z
Valid To: 2020-05-06T16:25:24.000Z

Thumbprint: 72:FE:95:F2:90:7C:86:24:D9:4E:12:EC:FB:10:38:7A:DA:EC:00:3A
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:25:19.000Z
Valid To: 2020-05-06T16:25:19.000Z

Thumbprint: FA:AD:FD:91:AD:E4:F1:00:EC:4A:D4:73:81:DB:B2:D1:20:35:DB:F2
Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84
Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84
Subject Alternate Name: localhost,127.0.0.1
PublicKey Algorithm: RSA
Valid From: 2015-05-07T16:24:45.000Z
Valid To: 2020-05-06T16:24:45.000Z
```

Sample Contents of Custom vRealize Operations Manager Certificates

For troubleshooting purposes, you can open a custom certificate file in a text editor and inspect its contents.

PEM Format Certificate Files

A typical PEM format certificate file resembles the following sample.

```
-----BEGIN CERTIFICATE-----
MIIFlDCCBLYgAwIBAgIKFYXYUwAAAAAAGTANBgkqhkiG9w0BAQ0FADBhMRMwEQYK
CZImiZPyLGQBGRYDY29tMRUwEwYKCZImiZPyLGQBGRYFdm13Y3MxGDAWBgoJkiaJ
<snip>
vKStQJNr7z2+pTy92M6FgJz3y+daL+9ddbaMnp9fVXjHBoDLGGaLOvyD+KJ8+xba
aGJfGf9ELXM=
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQE415ffX694rii1RmdRLJwL6sOWa+Wf70HRoLtx21kZzbXbUQN
mQhTRiidJ3Ro2gRbj/btSsI+OMUzotz5VRT/yeyoTC5l2uJEapld45RroUDHQwWJ
<snip>
DAN9hQus3832xMkAuVP/jt76dHDYyviyIYbmXzMalX7LZy1MCQVg4hCH0vLsHtLh
MlrOAsz62Eht/iB61AsVCCiN3gLrX7MKsYdxZcRVruGXSIh33ynA
-----END RSA PRIVATE KEY-----
```

```
-----BEGIN CERTIFICATE-----
MIIDnTCCAoWgAwIBAgIQY+j29InmdYNCs2cKlH4kPzANBgkqhkiG9w0BAQ0FADBh
MRMwEQYKCZImiZPyLGBGRYDY29tMRUwEwYKCZImiZPyLGBGRYFdm13Y3MxGDAW
<snip>
ukzUuqX7wEhc+QgJWgl41mWZBZ09gfsA9XuXBL0k17IpVHpEgwrrjQz8X68m4I99
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: 1e-WebServer-8dea65d4-c331-40f4-aa0b-205c3c323f62
Key Attributes
X509v3 Key Usage: 10
-----BEGIN PRIVATE KEY-----
MIICdwIBADANBgkqhkiG9w0BAQEFAASCAmEwgGJdAgEAAoGBAKHqyfc+qcQK4yxJ
om3PuB8dYZm34Qlt81GAAnBPYe3B4Q/0ba6PV8GtWG2svIpc1/eflwGHgTU3zJxR
gkKh7I3K5tGESn81ipyKtKpYebh+aBMqPKrNNUEKlr0M9sa3WSc0o3350tCc1ew
5ZkNYZ4BRUVVYm0HogeGhOthRn2fAgMBAAECGyABhPmGN3FSZKPDG6HJ1ARvT1BH
KAGVnBGHd0MOMABghFBnBKXa8LwDldgGBngloOakEXTftkIjdB+uwkU5P4aRrO7
vGuJUtRyRCU/4fjLBDuxQL/KpQfruAQaof9uWUwh5W9fEeW3g26fzVL8AFZnbXS0
7Z0AL1H3LNLd5rpOQJBANnI7vFu06bFxFVF+kq6ZOJFMx7x3K4VGxgg+PfFEBEPS
UJ2LuDH5/Rc63BaxFzM/q3B3Jhehvgw61mMyxU7QSSUCQQC+VDuW3XEWJjSiU6KD
gEGpCyJ5SBePbLSukljpGidKkDN1kLgbWVytCVkTAmuoAz33kMWfqIiNcqQbUgVV
UnpzAkB7d0CPO0deSsy8kMdTmKXLKf4qSF0x55epYK/5MZhBYuA1ENrR6mmjW8ke
TDNc6IGm9sVvrFBz2n9kKYpWThrJAKeAk5R69DtW0cbkLy5MqEzOHQauP36gDi1L
WMXPvUfzSYTQ5aM2rrY2/1FtSSkqUwfYh9sw8eDbqVpIV4rc6dDfcwJBALiIDPT0
tz86wySJNeOiUkQm36iXVF8AckPKT9TrbC3Ho7nC8OzL7gEl1ETa4Zc86Z3wpcGF
BHhEDMHaihyuVgI=
-----END PRIVATE KEY-----
Bag Attributes
localKeyID: 01 00 00 00
1.3.6.1.4.1.311.17.3.92: 00 04 00 00
1.3.6.1.4.1.311.17.3.20: 7F 95 38 07 CB 0C 99 DD 41 23 26 15 8B E8
D8 4B 0A C8 7D 93
```

```

friendlyName: cos-oc-vcops
1.3.6.1.4.1.311.17.3.71: 43 00 4F 00 53 00 2D 00 4F 00 43 00 2D 00
56 00 43 00 4D 00 35 00 37 00 31 00 2E 00 76 00 6D 00 77 00 61 00
72 00 65 00 2E 00 63 00 6F 00 6D 00 00 00
1.3.6.1.4.1.311.17.3.87: 00 00 00 00 00 00 00 00 02 00 00 00 20 00
00 00 02 00 00 00 6C 00 64 00 61 00 70 00 3A 00 00 00 7B 00 41 00
45 00 35 00 44 00 44 00 33 00 44 00 30 00 2D 00 36 00 45 00 37 00
30 00 2D 00 34 00 42 00 44 00 42 00 2D 00 39 00 43 00 34 00 31 00
2D 00 31 00 43 00 34 00 41 00 38 00 44 00 43 00 42 00 30 00 38 00
42 00 46 00 7D 00 00 00 70 00 61 00 2D 00 61 00 64 00 63 00 33 00
2E 00 76 00 6D 00 77 00 61 00 72 00 65 00 2E 00 63 00 6F 00 6D 00
5C 00 56 00 4D 00 77 00 61 00 72 00 65 00 20 00 43 00 41 00 00 00
31 00 32 00 33 00 33 00 30 00 00 00
subject=/CN=cos-oc-vcops.eng.vmware.com
issuer=/DC=com/DC=vmware/CN=VMware CA
-----BEGIN CERTIFICATE-----
MIIFWTCCBEGgAwIBAgIKSJGT5gACAAAwKjANBgkqhkiG9w0BAQUFADBMMwEQYK
CZImizPyLGQBGRYDY29tMRYwFAYKCZImizPyLGQBGRYGdm13YXJlMRIwEAYDVQQD
Ew1WTXdhcmUgQ0EwHhcNMTQwMjA1MTg1OTM2WhcNMTYwMjA1MTg1OTM2WjAmMSQw

```

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 4-224. 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 4-225. Certificate Options

Option	Description
Certificate Thumbprint	Unique alphanumeric string associated with the certificate
Issued By	Content associated with the issuer of the certificate, such as organization name and location
Issued To	Typically, content associated with the issuer, plus the certificate object Identifier (OID)
Expires	The date after which the certificate cannot be used for successful authentication

Add a Custom Certificate to vRealize Operations Manager

If you did not add your own SSL/TLS certificate when configuring the vRealize Operations Manager primary node, you can still add a certificate after vRealize Operations Manager is installed.

Prerequisites

- Create and configure the primary node.
- Verify that your certificate file meets the requirements for vRealize Operations Manager. See the *vRealize Operations Manager vApp Deployment and Configuration Guide* or *vRealize Operations Manager Installation and Configuration Guide for Linux and Windows*.

Procedure


- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at `https://node-FQDN-or-ip-address/admin`.
- 2 Log in with the admin user name and password.
- 3 At the upper right, click the SSL certificate icon.
- 4 In the certificate window, click **Install New Certificate**.
- 5 Click **Browse for certificate**.
- 6 Locate the certificate `.pem` file, and click **Open** to load the file in the Certificate Information text box.
- 7 Click **Install**.

Removing an Adapter Certificate


If you want to delete an old or expired certificate associated with an adapter, perform the following steps:

Procedure

- 1 In a Web browser, navigate to the vRealize Operations Manager administration interface at <https://node-FQDN-or-ip-address/ui>.
- 2 Log in with the administrator user name and password.
- 3 In the menu, click **Administration**, and in the left pane click **Management > Certificates**.
- 4 In the certificate window, select the certificate that has to be removed.
- 5 Click **Delete** to remove the certificate.
- 6 If the certificate is being used by the adapter, then the following message comes up:



A certificate can be configured for one or more adapters if it is the same destination system.

- 7 If you delete a certificate which is already being used by another adapter, the adapter fails to connect or start. As a workaround, perform the following steps:
 - a On the left pane, click **Solutions**.
 - b Select the particular adapter and click the Configure button  on the toolbar.
 - c Click **Test Connection**.
 - d A prompt comes up asking the user to import the associated certificate. Click **OK**.
 - e Restart the adapter from the **Solutions** page.

Upgrade Internal Certificates

The internal certificates for vRealize Operations Manager expire five years after its initial installation. Upgrade your internal certificates for vRealize Operations Manager 6.3 and later versions using the certificate renewal PAK file. After logging in, if you see a message like, "vRealize Operations Manager internal certificates will expire on mm/dd/yyyy. Please install a new certificate before the expiry date. For more details, see KB 71018." in the Quick Start page, you must upgrade your internal certificates for vRealize Operations Manager using the certificate renewal PAK file from the vRealize Operations Manager Administrator interface.

Prerequisites

- Obtain the PAK file for your cluster. See [Obtain the Software Update PAK File](#) for details.

Note The certificate renewal PAK is a standalone tool used only for the vRealize Operations Manager internal certificate renew.

- Bring your cluster offline before installing the PAK file to upgrade your internal certificates.

Note If your internal certificates have already expired, install the `vRealize_Operations_Manager_Enterprise_Certificate_Renewal_PAK` manually. For more information, see the following KB article [71018](#).

Procedure

- 1 Log into the vRealize Operations Manager administrator interface of your cluster at `https://master-node-FQDN-or-IP-address/admin`.
- 2 Take the vRealize Operations Manager cluster offline. For more information, see [vRealize Operations Manager Cluster Management](#).
- 3 Install the `vRealize_Operations_Manager_Enterprise_Certificate_Renewal_PAK` to upgrade your internal certificates. See, [Install a Software Update](#).

Note After the installation is complete, the administrator interface logs you out.

- 4 Log back into the vRealize Operations Manager administrator interface.
- 5 Bring the vRealize Operations Manager cluster back online.

After the cluster goes online, the upgrade is complete.

Modifying Global Settings

The global settings control the system settings for vRealize Operations Manager, including data retention and system timeout settings. You can modify one or more of the settings to monitor your environment better. These settings affect all your users.

The global settings do not affect metric interactions, color indicators, or other object management behaviors. These behaviors are configured in your policies.

Settings related to managing objects with vRealize Operations Manager are available on the **Inventory** page.

You can view tooltips for each option in the Edit Global Settings dialog box.

Global Settings Best Practices

Most of the settings pertain to how long vRealize Operations Manager retains collected and process data.

The default values are common retention periods. You might need to adjust the time periods based on your local policies or disk space.

List of Global Settings

The global settings determine how vRealize Operations Manager retains data, keeps connection sessions open, and other settings. These are system settings that affect all users. Some of these settings are not editable. Editable global settings have a hidden Edit icon next to their values. To see the icon, point to the global setting.

Table 4-226. 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	<p>Number of hours to retain objects that are deleted from an adapter data source or server before deleting them from vRealize Operations Manager.</p> <p>An object deleted from an adapter data source is identified by vRealize Operations Manager as not existing and vRealize Operations Manager can no longer collect data about the object. Whether vRealize Operations Manager identifies deleted objects as not existing depends on the adapter. This feature is not implemented in some adapters.</p> <p>For example, if the retention time is 360 hours and a virtual machine is deleted from a vCenter Server instance, the virtual machine remains as an object in vRealize Operations Manager for 15 days before it is deleted.</p> <p>This setting applies to objects deleted from the data source or server, not to any objects you delete from vRealize Operations Manager on the Inventory page.</p> <p>A value of -1 deletes objects immediately.</p> <p>You can define the number of hours per object type to retain objects that no longer exist and check for object type overrides. To add individual object types and set up their values, click the Object Deletion Scheduling icon. You can also edit or delete these object types.</p>
Deletion Scheduling Interval	24 hours	Determines the frequency to schedule deletion of resources. This setting works with the Deleted Objects setting to remove objects that no longer exist in the environment. vRealize Operations Manager transparently marks objects for removal that have not existed for the length of time specified under Deleted Objects. vRealize Operations Manager then removes the marked objects at the frequency specified under Deletion Scheduling Interval.
Object History	90 days	<p>Number of days to retain the history of the object configuration, relationship, and property data.</p> <p>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.</p> <p>The data is purged from the system after the specified number of days.</p>
Generated Reports Retention	Disabled	Number of months to retain generated reports. If disabled, all the generated reports will be retained.
Session Timeout	30 minutes	<p>If your connection to vRealize Operations Manager is idle for the specified amount of time, you are logged out of the application.</p> <p>You must provide credentials to log back in.</p>

Table 4-226. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Symptoms/Alerts	45 days	Number of days to retain canceled alerts and symptoms. The alerts and symptoms are either canceled by the system or by a user.
Time Series Data Retention	6 months	Number of months that you want to retain the collected and calculated metric data for the monitored objects. This setting is set to 6 months by default for 5 minutes interval data retention.
Additional Time Series Retention	36 months	The number of months that the roll-up data extends beyond the regular period. The roll-up data is available starting from the end of the regular period and until the end of the roll-up data retention period. If you specify 0 as the value, then this will effectively disable the Additional Time Series Data Retention time and only data specified in Time Series Retention is stored. This setting ensures that after 6 months of normal retention for 5 minutes, the seventh month data is rolled up into a one Hour roll up. You can set up this option up to 120 months for data roll ups.
Deleted Users	100 days	You can specify the number of days to keep custom content created by a user who has been removed from vRealize Operations Manager or by the automatic synchronization of LDAP. For example, the custom dashboards created by a user.
External Event Based Active Symptoms	disabled	The number of days to retain the external event-based active symptoms.
Maintain Relationship History		You can maintain a history of all the relationships of all the monitored objects in vRealize Operations Manager .
Dynamic Threshold Calculation	enabled	Determines whether to calculate normal levels of threshold violation for all objects. If the setting is disabled, the following area of vRealize Operations Manager does not work or are not displayed: <ul style="list-style-type: none"> ■ Alert symptom definitions based on dynamic thresholds will not work ■ Metric charts that display normal behavior are not present Disable this setting only if you have no alternative options for managing resource constraints for your vRealize Operations Manager system.
Cost Calculation		The host time at which cost calculations are run.
Customer Experience Improvement Program	enabled	Determines whether to participate in the Customer Experience Improvement Program by having vRealize Operations Manager send anonymous usage data to https://vmware.com .

Table 4-226. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Allow vCenter users to log in to individual vCenters using the vRealize Operations Manager UI		<p>Lists all the individual vCenter Server on the vRealize Operations Manager login page to allow users to use their individual vCenter Server credentials to log into vRealize Operations Manager.</p> <ul style="list-style-type: none"> ■ vCenter Server users can log in from vCenter Server clients. Enabled by default, this lists all the configured vCenter Servers in the vRealize Operations Manager login page drop-down.
Allow vCenter users to log in from vCenter clients	enabled	Allows vCenter Server users to log in from the vCenter Server clients.
Allow vCenter users to log in to all vCenters using the vRealize Operations Manager UI	enabled	<p>List all the vCenter Servers on the vRealize Operations Manager login page to allow users to use their vCenter Server credentials to log into vRealize Operations Manager.</p> <p>Allows vCenter Server users to log in vRealize Operations Manager UI using any vCenter Server credentials.</p> <p>Enabling this option adds all vCenter Server in the vRealize Operations Manager login page dropdown.</p>
System access URL		You can specify the URL that is used to access the system when a load balancer is used. The URL that you enter here is displayed in the outbound notifications and while sharing dashboards.
Automated Actions	enabled or disabled	Determines whether to allow vRealize Operations Manager to automate actions. When an alert triggered, the alert provides recommendations for remediation. You can automate an action to remediate an alert when the recommendation is the first priority for that alert. You enable actionable alerts in your policies.
Enable Standard Certification Validation		<p>This option enables certificate verification to Test Connection in the Create or Modify AI screen, using a standard verification flow.</p> <p>The option checks CA authority.</p> <ul style="list-style-type: none"> ■ Certificate Subject DN ■ Subject alternative name ■ Certificate validity period ■ Revocation list <p>This option also presents dialogs to user if one of those checks fail. It is up to the adapter implementation on how the adapter checks source certificate validity during a normal collection cycle. On a usual scenario, adapters just perform a thumb-print verification. However, in case this flag is enabled, Test connection validates certificates in full scale and accepts certificates that are matching all criteria without any user dialogs.</p>
Concurrent UI login sessions	enabled	Allows concurrent UI login sessions per user. Once changed, this setting affects the subsequent login sessions.

Table 4-226. Global Setting Default Values and Descriptions (continued)

Setting	Default Value	Description
Allow non-imported vIDM user access	enabled	Allows non-imported VMware Identity Manager users to be created automatically as read-only users upon first access. If disabled, only VMware Identity Manager imported users or users belonging to imported VMware Identity Manager groups will be granted access.
Currency		You can specify the currency unit that is used for all the cost calculations. You can select the type of currency from the list of currency types by clicking Choose Currency . From the Set Currency , select the required currency and confirm your action by clicking the check box, and set the currency.

Global Settings

To manage how vRealize Operations Manager retains data, keeps connection sessions open, and other settings, you can modify the values for the global settings. These system settings affect all users.

You can also choose to participate in the customer experience improvement program. For more information on accessing Global settings, see [Access Global Settings](#).

Access Global Settings

With global settings, you set times to delete objects, set timeouts, store historical data, use dynamic threshold and capacity calculations, and determine how vCenter Server users log in. For automated actions, you can select whether to allow actions to be triggered from alert recommendations automatically.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane click **Management > Global Settings**.
- 2 To edit the global settings, click the setting you want to edit.

Note Editable global settings have a hidden **Edit** icon next to their values. To see the icon, point to the global setting.

Table 4-227. Global Settings Options

Option	Description
Edit Global Settings	Click the global setting you want to edit to activate the edit mode and modify the setting values. To edit non-switchable settings, select a value and then click Save . To edit switchable settings, select a value and then click Enable or Disable to change the setting. Click Cancel to discard all changes and exit the edit mode.
Setting	Setting name.

Table 4-227. Global Settings Options (continued)

Option	Description
Value	Current value for the setting. To change the setting value, click Edit Global Settings .
Description	Information about the setting. Point to the setting to display additional information about the setting.

The Customer Experience Improvement Program

This product participates in VMware's Customer Experience Improvement Program (CEIP). The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. You can choose to join or leave the CEIP for vRealize Operations Manager at any time.

This product also uses a JavaScript operated by VMware's service provider Pendo.io. The JavaScript collects information on your interactions with the user interface such as clickstream data, page loads, limited browser, and device information. It helps VMware to understand how the product is used. This data is used to improve VMware products and services and design them better. For more information, see [VMware's Privacy Notices](#).

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

You can opt-out from such data collection by opting out of VMware's CEIP program. Additional controls are also provided to individual users in the user interface.

Join or Leave the Customer Experience Improvement Program for vRealize Operations Manager

You can join or leave the Customer Experience Improvement Program (CEIP) for vRealize Operations Manager at any time.

vRealize Operations Manager gives you the opportunity to join the Customer Experience Improvement Program (CEIP) when you initially install and configure the product. After installation, you can join or leave the CEIP by following these steps.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Management > Global Settings**.
- 2 From the toolbar, click the **Edit** icon.
- 3 Select or clear the **Customer Experience Improvement Program** option.
This option activates the program and sends data to www.vmware.com.
- 4 Click **OK**.

Managing Content

As a vRealize Operations Manager admin, you can take regular backups of your custom and the out-of-the-box content to manage your operational or regulatory needs. If there is a wrong edit or if the need to recover data arises, then you can use the recent backup to restore the content or import the content to a different setup. By taking regular backups, you can also upgrade the vRealize Operations Manager to the latest build without losing or overriding the custom content.

Note Any user with the "View Content Export/Import Page" permission can export the content. However, only a super admin has the privilege to export all the content, including the content owned by other users, for example, custom dashboards.

Creating a Backup

You can create regular backups of your custom and the out-of-the-box content in vRealize Operations Manager . You can use this backup to restore your content or export the content while setting up another environment.

You can take a backup of the following content types available in vRealize Operations Manager .

- Alert Definitions
- Custom Compliance Benchmarks
- Custom Groups
- Dashboards
- Metric Configurations
- Notification Rules
- Policies
- Recommendations
- Report Templates
- Super Metrics
- Symptom Definitions
- Views

Procedure

- 1 In the **Administration** page, click **Management > Content Management**.
- 2 In the **Export Content** tab, click **Generate Export Content** to create a backup.

The system compresses the content into one ZIP file.

- 3 Click the **Download ZIP file** link to download the backup content.

You can use the downloaded content to restore your content or export it to a different setup.

Importing Content

You can take regular backups of your custom and the out-of-the-box content and import it to a different environment.

Prerequisites

- Ensure that you have downloaded the backup ZIP file. For details, see [Creating a Backup](#).
- Ensure that all the users who own the custom dashboards are present in the destination setup so that the custom dashboards are assigned to the respective owners when the content is imported. Otherwise, the dashboards of the owners who are not present in the destination setup will be skipped while importing the content.

Procedure

- 1 In the **Administration** page, click **Management > Content Management**.
- 2 Click the **Import Content** tab and then, click **Browse** to select the downloaded ZIP file.
- 3 If there is a conflict while importing the content, you can select to either **Override existing content** or **Skip item(s)**.

The details of the overridden or the skipped content are displayed only during the import and right after the import is completed. You can view this information under the **Results** section in the same page.

- 4 Click **Import Content**.

After the import is completed, the content is available in the destination setup.

Transfer Ownership of Dashboards and Report Schedules

When a user is deleted from vRealize Operations Manager, the report schedules and dashboards created by the user are stored as orphaned content. As an admin user, you can transfer ownership of dashboards and report schedules created by deleted users.

From Where You Can Transfer Ownership of Dashboards and Report Schedules

In the menu, click **Administration**. From the left pane, select **Management > Orphaned Content**.

Orphaned Content Page

You can view a list of deleted users from the **Deleted Users** panel in the left pane of the **Orphaned Content** page. Based on your selection in the **Deleted Users** panel, the dashboards, and report schedules for the deleted user are displayed under the **Dashboard** and **Report Schedules** tabs in the **Orphaned Content** page.

As an admin user, you can take ownership, assign ownership, or discard orphaned dashboards and report schedules, from the **Actions** menu in the **Dashboards** and **Report Schedules** tabs. Enter the name or part of the name of a dashboard or report schedule in the **Filter** option and click **Enter**. The relevant dashboard or report schedule is displayed.

Table 4-228. Actions Menu Options

Actions	Options
Take Ownership	You can take ownership of the selected dashboards or report schedules.
Assign Ownership	You can assign a new owner for the selected dashboards or report schedules. You can select a target user from the Transfer Dashboards/Report Schedule dialog box.
Discard	You can permanently delete the dashboards or report schedules.

vRealize Operations Manager Logs for Product UI

How vRealize Operations Manager Logs Work

For troubleshooting in the product UI, the product provides an expandable tree of vRealize Operations Manager log files that you can browse and load for review. You can also edit the log file folders, limit the retained log size, and set logging levels.

vRealize Operations Manager logs are categorized by cluster node, and log type. All logs are in the UTC formatted date and time. The logging format is as follows:

```
Date/Time+0000, LEVEL, [THREAD/IP Address], [Specific Fields], CLASS - MESSAGE
```

If you have configured a timezone for the vRealize Operations Manager VM, the system logs will be in that timezone. The vRealize Operations Manager logs will remain in UTC.

Where You Find vRealize Operations Manager Logs

In the menu, click **Administration**, and in the left pane click **Support > Logs**.

Log Viewer Options

Use the toolbar options to control the tree of items and the viewer.

- 1 Click **Node** and select any component that is listed under the node.
- 2 Click the gear icon, enter the logging levels and log size.
- 3 Click **OK**.

Note Not all components have relevant syslog information. Therefore, not all nodes have the configuration option enabled.

Figure 4-3. Logs




Figure 4-4. Log Options




Table 4-229. Log Viewer Toolbar Options

Option	Description
Group By	Organizes the tree by cluster node or log type.
Collapse All	Closes the view of the tree to show only the high-level folders.
Edit Properties	For the selected folder, you can limit the log size and set logging levels.
Delete Selected File	Deletes the log file.
Starting Line	Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end.
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 .

Table 4-229. Log Viewer Toolbar Options (continued)

Option	Description
Text to Find	<p>Enter the specific text that you want to search in the logs. Add the following filters for search, if required:</p> <ul style="list-style-type: none"> ■ Case Sensitive ■ Regular Expression <p>You can perform the search at various levels:</p> <ul style="list-style-type: none"> ■ 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. <p>The last modified time for any file is found by placing the pointer on the file in the tree.</p>
Set Timerange	<p>If you specify a time range, the logs for that particular time range are shown in the search results.</p>
Word Wrap	<p>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.</p>

Create a vRealize Operations Manager Support Bundle

You create a vRealize Operations Manager support bundle to gather log and configuration files for analysis when troubleshooting a vRealize Operations Manager issue.

When you create a support bundle, vRealize Operations Manager gathers files from cluster nodes into ZIP files for convenience.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Support > Support Bundles**.
- 2 From the toolbar, click the **Create a Support Bundle** icon.
- 3 Select the option to create a **Light** or **Full support bundle**.
- 4 Select the cluster nodes that need to be evaluated for support.
Only logs from the selected nodes are included in the support bundle.
- 5 Click **OK**, and click **OK** to confirm support bundle creation.
Depending on the size of the logs and number of nodes, it might take time for vRealize Operations Manager to create the support bundle.

What to do next

Use the toolbar to download the support bundle ZIP files for analysis. For security, vRealize Operations Manager prompts you for credentials when you download a support bundle.

You can review the log files for error messages or, if you need troubleshooting assistance, send the diagnostic data to VMware Technical Support. When you resolve or close the issue, use the toolbar to delete the outdated support bundle to save disk space.

vRealize Operations Manager Support Bundles

vRealize Operations Manager support bundles contain log and configuration files that help troubleshoot a vRealize Operations Manager issue.

How Support Bundles Work

Support bundles require that you select nodes or the entire cluster, and the level of logging that you want to collect. After vRealize Operations Manager creates the support bundle, you download it in ZIP format for analysis.

Where You Find Support Bundles

In the menu, click **Administration**, and then in the left pane, select **Support > Support Bundles**.

Support Bundle Options

The options include toolbar and data grid options.

You can click **Add** or click the **Horizontal Ellipses** to delete, download, or reload support bundles.

Table 4-230. Support Bundle Toolbar Options

Option	Description
Add	Open a dialog box that guides you through the process of creating a support bundle.
Delete	Remove the selected support bundle.
Download	Download the support bundle in ZIP format.
Reload Support Bundles	Refresh the list of support bundles.

Use the data grid options to view item details.

Table 4-231. Support Bundle Data Grid Options

Option	Description
Bundle	System-generated identifier for the support bundle.
Bundle Type	<ul style="list-style-type: none"> ■ 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 4-232. Dynamic Threshold Options

Option	Description
Start	Run the dynamic threshold calculation process now, outside of its normal schedule.
Stop	Stop the dynamic threshold calculation currently in progress.
Calculation progress	Percentage completion of the current dynamic threshold calculation.
Calculation times and Count	Timestamps and metric counts associated with the last dynamic threshold calculation, and the time for the next scheduled calculation.

vRealize Operations Manager Adapter Redescribe

When vRealize Operations Manager redescribes an adapter, vRealize Operations Manager finds the adapter files, gathers information about the abilities of the adapter, and updates the user interface with information about the adapter.

How Adapter Redescribe Works

After installing or updating an adapter, capture the adapter information by having vRealize Operations Manager redescribe its adapters.

Where You Find Adapter Redescribe

In the menu, click **Administration**, and then in left pane, click **Support > Redescribe**.

Adapter Redescribe Options

The feature includes an option to start the adapter describe process.

Table 4-233. Adapter Redescribe Options

Option	Description
Redescribe	Start the adapter describe process.

vRealize Operations Manager provides adapter-specific details from the redescribe process.

Table 4-234. 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 <code>describe.xml</code> against which the last redescribe process ran.
Adapter Version	Version of the adapter against which the last redescribe process ran.
Message	Additional details about the last redescribe process.

Customizing Icons

Every object or adapter in your environment has an icon representation. You can customize how the icon appears.

vRealize Operations Manager assigns a default icon to each object type and adapter type. Taken collectively, object types and adapter types are known as objects in your environment. Icons represent objects in the UI and help you to identify the type of object. For example, in the Topology Graph widget on a dashboard, labeled icons show how objects are connected to one other. You can quickly identify the type of object from the icon.

If you want to differentiate objects, you can change the icon. For example, a virtual machine icon is generic. If you want to pictorially distinguish the data that a vSphere virtual machine provides from the data that a Hypervisor virtual machine provides, you can assign a different icon to each.

Customize an Object Type Icon

You can use the default icons that vRealize Operations Manager provides, or you can upload your own graphics file for an object type. When you change an icon, your changes take effect for all users.

Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons**.
- 2 Click the **Object Type Icons** tab.
- 3 Assign the Object Type icon.
 - a Select the object type in the list with the icon to change.
By default, object types for all adapter types are listed. To limit the selection to the object types that are valid for a single adapter type, select the adapter type from the drop-down menu.
 - b Click the **Upload** icon.
 - c Browse to and select the file to use and click **Done**.
- 4 (Optional) To return to the default icon, select the object type and click the **Assign Default Icons** icon.

The original default icon appears.

Object Type Icons Tab

vRealize Operations Manager obtains data from different sources. Data sources are classified by the type of object or object type. In UI locations where metric data appears for objects, vRealize Operations Manager includes an icon to show the object type. To graphically distinguish the different types of objects, you can customize the icon.

Where You Customize Object Type Icons

In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons > Object Type Icons**.

Table 4-235. 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. <ul style="list-style-type: none"> ■ Upload uploads a PNG file to uniquely identify the object type. ■ Assign Default icons returns the selection to the original icon.
Search	Search for objects with a particular name to narrow the selection of object types displayed.
Object Type	Name of the type of object.
Icon	Pictorial representation of the type of object.

Customize an Adapter Type Icon

You can use the default icons that vRealize Operations Manager provides, or you can upload your own graphics file for an adapter type. When you change an icon, your changes take effect for all users.

Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

Procedure

- 1 In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons**.
- 2 Click the **Adapter Type Icons** tab.
- 3 Assign the Adapter Type icon.
 - a Select the adapter type in the list with the icon to change.
 - b Click the **Upload** icon.
 - c Browse to and select the file to use and click **Done**.
- 4 (Optional) To return to the default icon, select the adapter type and click the **Assign Default Icons** icon.

The original default icon appears.

Adapter Type Icons Tab

Adapters collect and provide data to vRealize Operations Manager. Adapters are classified by the type of adapter or adapter kind. To graphically distinguish the different types of adapters, you can customize the icon.

Where You Customize Adapter Type Icons

In the menu, click **Administration**, and then in the left pane, click **Configuration > Icons > Adapter Type Icons**.

Table 4-236. Adapter Type Icons Options

Option	Description
Toolbar options	<p>Manages the selected icon.</p> <ul style="list-style-type: none"> ■ 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.

- Enable or disable high availability (HA).
- Upload and install vRealize Operations Manager software update PAK files.
- The product interface is inaccessible, and you must correct the problem by bringing nodes online, or by restarting nodes or the cluster.
- vRealize Operations Manager needs to be restarted for any reason.

There is some overlap between the administration interface and product interface in terms of access to logs, support bundles, and some of the node maintenance activities that do not involve restarting the cluster, such as adding nodes.

vRealize Operations Manager Cluster Management

vRealize Operations Manager includes a central page where you can monitor and manage the nodes in your vRealize Operations Manager cluster and the adapters that are installed on the nodes.

How Cluster Management Works

You can view and change the online or offline state of the overall vRealize Operations Manager cluster or the individual nodes. In addition, you can enable or disable high availability (HA) and view statistics related to the adapters that are installed on the nodes.

Where You Find Cluster Management

Log in to the vRealize Operations Manager administration interface at <https://master-node-name-or-ip-address/admin>.

Cluster Management Options

The options include cluster-level monitoring and management features.

Table 4-237. Initial Setup Status Details

Option	Description
Cluster Status	<p>Displays the online, offline, or unknown state of the vRealize Operations Manager cluster and provides an option to take the cluster online or offline.</p> <p>If a cluster fails to go offline, click the Force Take Offline button to take the cluster offline.</p> <p>Note The Force Take Offline button appears only when the Bring Cluster offline operation fails.</p> <p>You can select to display the reason for taking the cluster offline. Select the Show reason on maintenance page check box in the Take Cluster Offline dialog box. When you log in to vRealize Operations Manager when the cluster is offline, the reason for taking the cluster offline is displayed.</p>
High Availability	Indicates whether HA is enabled, disabled, or degraded and provides an option to change that setting.
Continuous Availability	Indicates whether CA is enabled, disabled, or degraded and provides an option to change that setting.

vRealize Operations Manager provides node-level information as well as a toolbar for taking nodes online or offline.

Table 4-238. Nodes in the vRealize Operations Manager Cluster

Option	Description
Generate Passphrase	Generate a passphrase that can be used instead of the administrator credentials to add a node to this cluster.
Take Node Online/Offline	You can select the required node and bring it online or offline. You are required to understand the risk involved and provide a valid reason for the action performed when you bring a node online or offline.

Table 4-238. Nodes in the vRealize Operations Manager Cluster (continued)

Option	Description
Reload Nodes	You can fetch data from the nodes.
Shrink Cluster	<p>This option provides a mechanism to remove a node without having to lose any data. The shrink cluster removes nodes by migrating data from one node to any other node.</p> <p>All the historical data is either moved to the primary node or any other node, which has sufficient disk space.</p> <p>If HA is enabled and you have selected the replica node for removal, then you are asked to select another replica node. vRealize Operations Manager provides a list of nodes that be a possible candidate to become a replica node.</p> <p>vRealize Operations Manager stops collecting data from the removed nodes. However, the data that is available in the removed node is migrated to an existing node. Once the migration is complete, then the removed nodes are deleted with the cluster state as offline.</p> <p>For remote collectors, if any adapters are on the collectors of the removed nodes, then such nodes are migrated as well.</p> <p>Note vRealize Operations Manager cannot move pinned adapters. The adapter instances which were pinned on removed nodes do not move to another collector automatically. You must change the collector before starting the shrink cluster process.</p>

Table 4-239. Nodes in the vRealize Operations Manager Cluster

Option	Description
Node Name	<p>Machine name of the node.</p> <p>The node that you are logged into displays a dot next to the name.</p>
Node Address	Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes may use DHCP or static IP.
Cluster Role	Type of vRealize Operations Manager node: primary, data, replica, or remote collector.
State	Powered on, powered off, unknown, or other condition of the node.
Status	Online, offline, unknown, or other condition of the node.
Objects	Total environment objects that the node currently monitors.
Metrics	Total metrics that the node has collected since being added to the cluster.

Table 4-239. Nodes in the vRealize Operations Manager Cluster (continued)

Option	Description
Build	vRealize Operations Manager software build number installed on the node.
Version	vRealize Operations Manager software version installed on the node.
Deployment Type	Type of machine on which the node is running: vApp
SSH Status	Enable or disable the SSH Status.

In addition, there are adapter statistics for the selected node.

Table 4-240. 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 4-241. 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 4-242. 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 4-243. Support Bundle Data Grid Options

Option	Description
Bundle	System-generated identifier for the support bundle.
Bundle Type	<ul style="list-style-type: none"> ■ Light. Include 24 hours of logs. ■ Full. Include all available logs and configuration files.
Date and Time Created	Time when support bundle creation began.
Status	Progress of support bundle creation.

Update the Reference Database for vRealize Operations Manager

You can update the reference database to have the most updated version of the reference library. The reference database supplies default values for cost calculations.

Procedure

- 1 In the menu, click **Administration** and in the left pane click **Support > Cost Reference Database**.

The existing version of the reference database along with the date is displayed.

- 2 Click **Download Here**.

The latest version of the reference database is downloaded to the default location.

- 3 Click **Upload Reference Database** and select the reference database from the default download location.

Results

Note that the updated reference library values are reflected in the cost drivers only after the cost calculation process runs as per the schedule.

Enable FIPS - Admin UI

You can enable Federal Information Processing Standards (FIPS) for vRealize Operations Manager to make your environment FIPS compliant.

You can enable FIPS in the vRealize Operations Manager cluster at the time of installation or after vRealize Operations Manager is up and running. Adding FIPS at installation is less intrusive because the cluster has not yet started.

If the cluster is running, to enable FIPS, you must take the cluster offline. For more information, see [vRealize Operations Manager Cluster Management](#).

- 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 Click **Administrator Settings**.

Note The **Enable FIPS** button is disabled when the cluster is running.

- 5 Click **Enable FIPS** after you take your cluster offline.

Note Once you enable FIPS, you cannot disable the FIPS mode in the current setup. To revert to a FIPS disabled setup, you must re-deploy vRealize Operations Manager.

- 6 In the **Are you sure you want to enable FIPS** dialog box, read the note and provide your consent for enabling FIPS and then click **Yes**.

Note Once you enable FIPS, the cluster restarts and is not be available during this time. The cluster nodes are rebooted and once the cluster is online, all the nodes are FIPS enabled.

Configuring and Using Workload Optimization

Workload Optimization provides for moving virtual compute resources and their file systems dynamically across datastore clusters within a data center or custom data center.

Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.

Workload Optimization further enables you potentially to automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention automatically runs an action, a data center performs at optimum.

vRealize Automation Integration

When you add an instance to a vRealize Automation adapter or solution pack as well as to a vCenter Server adapter instance that is connected to the vRealize Automation server, using vRealize Automation-managed resources, vRealize Operations Manager automatically adds a custom data center for the vCenter Server, using vRealize Automation-managed resources.

On the vRealize Operations Manager side, to get the day2 chain configured, you must make the following initial configurations:

- 1 In vCenter Server, **Administration -> Solutions** and then add the VMware vSphere adapter instance for the vCenter Server that is configured as an endpoint in vRealize Automation Server.
- 2 In vCenter Server, **Administration -> Solutions** and then add the VMware vRealize Automation adapter instance for the server that will appear in the vRealize Operations Manager and vRealize Automation integration day2 chain.

vRealize Operations Manager can manage workload placement and optimization for the custom data centers that reside in vRealize Automation-managed clusters.

However, vRealize Operations Manager is not permitted to set tag policies for the custom data center. (At the Workload Optimization screen, the Business Intent window is not operational for vRealize Automation custom data centers.) When rebalancing a vRealize Automation custom data center, vRealize Operations Manager uses all applicable policies and placement principles from both systems: vRealize Automation and vRealize Operations Manager. For complete information on creating and managing vRealize Automation custom data centers that are managed by vRealize Operations Manager, see the vRealize Automation documentation.

Configuring Workload Optimization

Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

- 1 Configure the Workload Automation Details. See [Workload Automation Details](#).
- 2 If you do not use the AUTOMATE function in the Optimization Recommendation pane at the Workload Automation screen, configure the two Workload Optimization alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the actions calculated by Workload Optimization are run automatically. See [Configuring Workload Optimization Alerts](#)

Prerequisites

Workload Optimization acts on objects associated with the VMware vSphere Solution that connects vRealize Operations Manager to one or more vCenter Server instances. The virtual objects in this environment include a vCenter Server, data centers and custom data centers, cluster compute and storage resources, host systems, and virtual machines. Specific requirements:

- A vCenter Adapter configured with the actions enabled for each vCenter Server instance.
- A vCenter Server instance with at least two datastore clusters with sDRS enabled and fully automated.
- Any non-datastore clusters must have DRS enabled and fully automated
- Storage vMotion must be set to ON at Workload Automation Details. The default is On.
- You must have permission to access all objects in the environment.

Design Considerations

The following rules constrain the possible computer and storage resource moves that can be performed.

Note When vRealize Operations Manager suggests that you optimize clusters in a data center, the system does not guarantee it can run an optimization action. vRealize Operations Manager analytics can determine that optimization is desirable and can create a rebalancing plan. However, the system cannot automatically identify all the architectural constraints that may be present. Such constraints may prevent an optimization action, or cause an action in progress to fail.

- Moving compute and storage resources is allowed only within, not across data centers or custom data centers.
- Storage resources cannot be moved across non-datastore clusters. Storage can move only across datastore clusters that have sDRS fully automated.
- Compute-resource-only moves are permitted through shared storage.
- Virtual machines defined with affinity rules or anti-affinity rules are not to be moved.
- Virtual machines cannot be moved when residing on a local datastore, unless a storage swap exists on the local datastore.
- Virtual machines cannot be moved if they have data residing across multiple datastore clusters. Compute-only moves with similar shared storage are not permitted.
- A virtual machine cannot have data that resides across different storage types. For example, if a virtual machine has a VM disk on a datastore and a second VM disk on a datastore cluster, the virtual machine does not move, even when the datastore is shared with the destination or has swap on it.
- A virtual machine can use RDM so long as the destination datastore cluster can access the RDM LUN.
- A virtual machine can implement VM disks on multiple datastores inside a single datastore cluster.
- Workload Optimization may suggest moving virtual machines that are protected by vSphere Replication or Array Based Replication. You must ensure that all the clusters within a selected data center or custom data center have replication available. You can set up DRS affinity rules on virtual machines that you do not want moving across clusters.

Business Intent Workspace

You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. vRealize Operations Manager can be configured to leverage tags to define business-related placement constraints: VMs can only be placed on hosts/clusters with matching tags.

Where You Find Business Intent

From the Home page, click the chevron next to Optimize Performance on the left. Click Workload Optimization, select a data center or custom data center from the top row, and click **Edit** in the Business Intent window.

To edit Business Intent values, you must have privileges for Administration -> Configuration -> Workload Placement Settings -> Edit.

Establishing Business Intent

Tags are implemented in vCenter Server as *key:value* labels that enable operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

To specify tags considered for placement, first select the radio button for the type of object you want to associate with VMs in this business intent session: Clusters or Hosts.

The system provides several suggested categories. These categories are only suggestions. You must specify the actual categories in vCenter Server after you expand the section for a suggested category. For example, in section "Tier", you can specify the actual vCenter Server tag category that represents tier semantics, for instance, "service level".

- Operating System
- Environment
- Tier
- Network
- Other

Any actual categories you specify must first be created in vCenter Server.

Then you can associate tagged VMs with clusters or hosts, based on the rules for each type of tagging.

- 1 Click the chevron to the left of the first suggested category. A **tag category** field appears.
- 2 Click the drop-down menu indicator and choose a category from the list defined in vCenter Server.
- 3 Click the drop-down menu indicator in the Tag Name (Optional) field and choose a tag name from the list defined in vCenter Server.
- 4 Click **Include Tag**. All VMs with that tag are associated with the category.

Rules for Host-Based Placement

To set host level placement constraints, vRealize Operations Manager automatically creates and manages DRS rules. All conflicting user-created DRS rules are DISABLED.

These rules include the following:

- Any VM-VM affinity and anti-affinity rules.
- Any VM-Host affinity and anti-affinity rules.

You must check the selection box next to the statement, "I understand that vRealize Operations will disable all my current and future DRS rules".

Configuring Workload Optimization Alerts

vRealize Operations Manager provides two preconfigured alerts designed to work with the Workload Optimization feature. You must take additional action in the Policies area to turn on the alerts and automate them so that predetermined actions are run when the alerts fire.

The following preconfigured alerts are designed to work with the Workload Optimization feature:

- Data center performance can potentially be optimized in one or more clusters.
- Custom data center performance can potentially be optimized in one or more clusters.

The preconfigured alerts fire only if the AUTOMATE function is not turned on at the Workload Optimization screen. (**Home -> Optimize Performance -> Workload Optimization**).

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI pages and manage vCenter Server objects.

Procedure

- 1 Select **Administration** from the menu, then **Policies** from the left pane.
- 2 Click **Policy Library** and select the policy that includes settings for the relevant data centers and custom data centers, for example, **vSphere Solution's Default Policy**.
- 3 Click the **Vertical Ellipses** and then **Edit**.
- 4 Click #6 on the lower left, Alert/Symptom Definitions.
- 5 Search on "can potentially be optimized" to locate the two alerts you want.
- 6 The alerts are turned ON by default/inheritance (see the State column).
- 7 The alerts are not automated by default/inheritance (see the Automate column). To automate the alerts, click the menu symbol to the right of the inherited value and select the green check mark.

Results

Workload Optimization is fully automated for your environment.

What to do next

To confirm that actions are taken automatically, monitor rebalance activity at the Workload Optimization screen.

Using Workload Optimization

Use the Workload Optimization UI pages to monitor optimizing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and run actions directly.

vRealize Operations Manager monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use optimization functions to distribute a workload differently in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects of interest.

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see [Chapter 6 Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager](#).

The following examples demonstrate the primary ways you can use Workload Optimization to keep your data centers balanced and performing their best.

Example: Run Workload Optimization

As a virtual infrastructure administrator or other IT professional, you use Workload Optimization functions to identify points of resource contention or imbalance. In this example, you manually run an optimization action to consolidate demand.

When you log into vRealize Operations Manager, you see the Quick Start page. In the left-most column, Optimize Performance, is the alert 3 DATA CENTERS REQUIRING OPTIMIZATION.

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

Procedure

- 1 Click **Workload Optimization** in the Optimize Performance column.

The Workload Optimization page appears. Data centers are grouped by Criticality, with the three troubled data centers appearing in a carousel across the top of the page: DC-Bangalore-18, DC-Bangalore-19, DC-Bangalore-20. A Not Optimized badge appears in the lower right corner of each graphic.

- 2 If no data center is preselected, select DC-Bangalore-18 from the carousel.

Comprehensive data about the state of the data center follows.

- 3 Based on the available data, you determine an optimization action is required.

CPU workloads can be consolidated such that a host in Cluster 3 can be freed up.

Table 4-244. Panes and Widgets

Pane	Contents
Workload Optimization	Status shows as Not Optimized. A system message says, "You can consolidate workloads to maximize usage and potentially free up 1 host." The message reflects that you have set policies to emphasize consolidation as a goal in optimization moves. The system is saying you can free up a host through consolidation.
Settings	The current policy is Consolidate. The system advises: Avoid Performance Issues, Consolidate Workloads.
Cluster Workloads	Cluster 1 CPU Workload is 16%. Cluster 2 CPU Workload is 29%. Cluster 3 CPU Workload is 14%. Cluster 4 CPU Workload is 22%.

- 4 Click **OPTIMIZE NOW** in the Workload Optimization pane.

The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.

- 5 If you are satisfied with the projected results of the optimization action, click **NEXT**.

The dialog box updates to show the planned moves.

- 6 If you need more information about the VMs which are included or excluded in the plan, click **Download Report** to see the optimization plan. You can review the reasons for incompatibilities and why some VMs were excluded from the plan.
- 7 Optional: If you want to know the total optimization potential of the move, assuming that there were no incompatibilities and all your VMs can be included in the optimization plan, click **Cancel**, and go to the Optimization Potential tab in the Workload Optimization page. Click **Calculate Optimization Potential** to see the total optimization potential of your data center.

- 8 Review the optimization moves, then click **BEGIN ACTION**.

The system runs the compute and storage resource moves.

Results

The optimization action moved compute and storage resources from some clusters to other clusters in the data center, and so freed up a host on one cluster.

Note The Workload Optimization page refreshes every five minutes. Depending on when you run an optimization action, the system might not reflect the result for up to five minutes, or longer when longer-running actions extend the processing time.

What to do next

To confirm that your optimization action was completed, go to the Recent Tasks page by selecting **Administration** on the top menu, and clicking **History > Recent Tasks** in the left pane. In the Recent Tasks page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

Note Sometimes an optimizing action may be suggested, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that suggested optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation might incur stress in the future, then consolidation is not suggested.

Example: Schedule a Repeating Optimization Action

As a virtual infrastructure administrator or other IT professional, you determine that compute and storage resources in a given data center are volatile and a regularly scheduled optimization action can address the problem.

vRealize Operations Manager monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization page. Depending on what appears, you may determine that you must schedule optimization functions to distribute a workload more evenly in a data center or custom data center.

Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

Procedure

- 1 From the Home screen, click **Optimize Performance > Workload Optimization** in the left pane.
- 2 From the carousel of data centers across the top of the page, select a data center for which you want to schedule repeated optimization actions.
- 3 In the Workload Optimization pane, click **SCHEDULE**.
- 4 Give the schedule a name and choose a time zone.
- 5 Determine how often you want to repeat the optimization action and click the relevant **radio button** under Recurrence.

Depending on your selection under Recurrence, additional options appear to the right. In this instance, you choose to repeat the optimization daily.

- 6 Leave the current date and time.
- 7 Select the **Repeat every day** radio button.
- 8 Select the **Expire after** radio button and tick the counter up to 6.
- 9 Click **Save**.

Results

The optimization action repeats for six days, then stops.

At the Workload Optimization page, the Scheduled button appears in the upper right of the Workload Optimization pane if optimization actions are scheduled for the selected data center. If you want to edit or delete a schedule, click the **Scheduled** button. The Optimization Schedules page appears, where you can perform those actions.

Note If you schedule a number of optimization actions close together, and the optimization plans of two or more actions include overlapping functions, that is, they impact the same set of resources, the system shifts the actions into a queue. As a result, some actions may complete later than expected, with longer running actions and other potential system constraints extending the lag time. Optimization actions that do not overlap can run concurrently.

What to do next

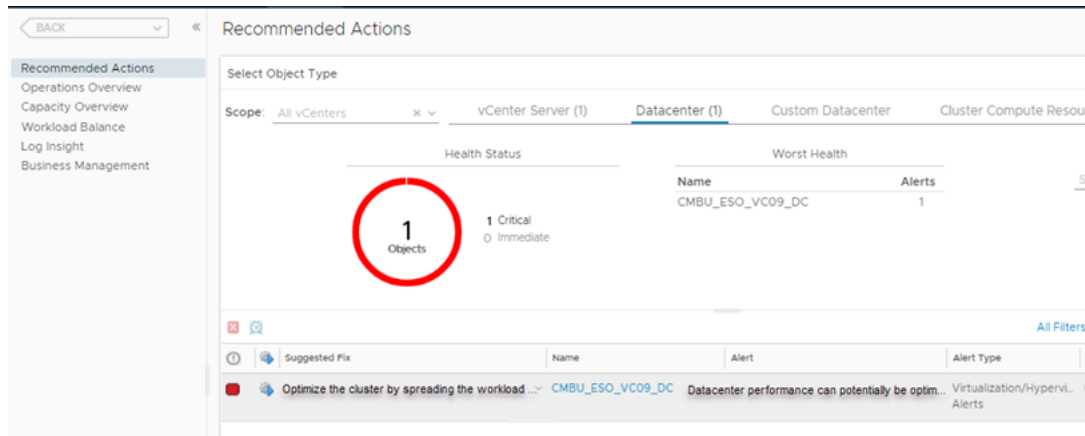
To confirm that your optimization action was finished, go to the Recent Tasks screen by selecting **Administration** on the top menu, and clicking **History > Recent Task** in the left pane. In the Recent Tasks screen, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, filter on Event Source and enter the name of the scheduled optimization plan.

Note Because real-time data center resource contention is dynamic, the system calculates a new optimization plan each time the scheduled optimization action starts, but before it runs. The system does not run the action if the system determines that the data center container is balanced at this moment. On the Recent Tasks page, the name of the affected data center appears in the Object Name column, and the Message “The optimization of the selected container cannot be improved” appears under Details. Another possibility is that a scheduled optimization plan is attempted, but does not go forward. In this event - which is not the same as a “failed” action - the name of the affected data center also appears in the Object Name column.

Example: Run Workload Optimization from Recommended Actions

From the Home screen, click **Recommendations** under Optimize Performance - first column on the left. The Recommended Actions screen appears, with data center and custom data center errors highlighted. If a suggested optimization action is available, it appears in the bottom third of the screen, with details.

To run the action, click the blue **Run Action** arrow.



Prerequisites

Ensure that you have all required permissions for accessing the Workload Optimization UI and managing vCenter Server objects.

Results

The system runs the proposed rebalancing action.

What to do next

The Workload Optimization screen appears, where you can review the results of the rebalancing actions. Additional information is available at the Recent Tasks page: in the menu, select **Administration**, then click **History > Recent Tasks** in the left pane. Choose the **Event Source** filter and enter part of the alert name, then search. If the action succeeded, the Event Source column shows Alert: *<alert name>*.

Workload Optimization Page

Workload Optimization enables you to optimize virtual machines and storage across datastore clusters to reduce resource contention and maintain optimum system performance.

Where You Find Workload Optimization

From the Home screen, select **Workload Optimization** under Optimize Performance in the left pane. From the Quick Start screen, select **Workload Optimization** in the left-most column.

Workload Optimization Page Options

In the Workload Optimization page, you see a list of data centers in a carousel, listed under three categories:

- Critical
- Normal
- Unknown

After you select a data center, you see the **ALL DATACENTERS** button on the upper right. Click **ALL DATACENTERS** when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers.

Table 4-245. Workload Optimization Page Options

Option	Description
View:	Filter results to include data centers, custom data centers, vRA-managed custom data centers, or all three. (Option appears if you select ALL DATACENTERS on the upper right.)
Group By:	Filter results by criticality (most out of balance data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. (Option appears if you select ALL DATACENTERS on the upper right.)
Sort By:	Options (Options appear if you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Alarm clock graphic - list data centers/custom data centers by time remaining. ■ Dollar sign - list data centers/custom data centers by potential cost savings with capacity optimization. ■ Scales graphic - Optimized.
Select data center or ADD NEW CUSTOM DATACENTER	Options (Options appear if you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object. ■ Select ADD NEW CUSTOM DATACENTER to display a screen that enables you to define a custom data center.

Data Center Options

After you select a data center from the carousel, you see the following information and options.

Note If you point your cursor to the lower right of a data center graphic, a tooltip may appear to let you know that the data center is using automated optimization.

Optimization Status Tab

Appears when you select a data center or custom data center from the top of the screen.

Table 4-246. The Optimization Recommendation Card

Option	Description
Status	<ul style="list-style-type: none"> ■ Optimized - indicates that workloads are optimized based on the settings you entered in the neighboring Operational Intent window, with no tag violations based on the settings you entered in the Business Intent window. ■ Not Optimized - indicates that one of the following conditions is true: workloads are not optimized based on the settings you entered in the neighboring Operational Intent window AND/OR there are tag violations based on the settings you entered in the Business Intent window. In the event of tag violations, the offending tags are listed.
OPTIMIZE NOW	Runs optimizing actions based on the settings you entered in your Operational and Business Intent settings.
SCHEDULE	Displays a dialog box enabling you to schedule one or more optimization actions. If schedules are currently set for data center or custom data center optimization, a check mark appears next to the data center or custom data center name.
AUTOMATE	<p>Continually seeks optimizing opportunities for data center or custom data center, based on the settings in the neighboring Operational Intent window or Business Intent windows. Scheduled optimizations are turned off while automatic optimization is on. Also, automated alerts are not operational when automatic optimization is on. Once you confirm automation, the system displays message, for example, 1) "Workload Optimization is looking for opportunities to automate," 2) "Your workloads are optimized according to your settings." or 3) "No eligible moves were found within the max number of compatibility checks allowed."</p> <p>Note To initiate Automation, you must have privileges for Environment -> Action -> Schedule Optimize Container.</p>
TURN OFF AUTOMATION	Stops automatic optimization. Any scheduled optimizations come back online.

Note Sometimes an optimizing action may be recommended, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that recommended optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation can incur stress in the future, then consolidation is not recommended.

Table 4-247. The Operational Intent Card

Option	Description
Utilization Objective	Indicates the main attribute of your current automation policy settings. Values are moderate, consolidate, or balance.
Edit	Displays the Workload Automation Policy Settings, where you can adjust settings for optimization and cluster headroom.

Table 4-248. The Business Intent Card

Option	Description
Intent	Allows you to define zones of infrastructure within cluster boundaries.
Edit	Displays a workspace where you can select criteria for placement of VMs.

Table 4-249. Details for Are your clusters meeting your utilization objective?

Option	Description
Are your clusters meeting your utilization objective?	<p>Displays a table which presents data in the following columns:</p> <ul style="list-style-type: none"> ■ Name ■ CPU Workload ■ Memory Workload ■ DRS Settings ■ Migration Threshold ■ Violated Tags ■ VM Name <p>Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. The violated tags shows which clusters or host groups are breaching the business intent. The VM Name column shows the name of the VMs and tag value due to which tag violation is happening.</p> <p>Provides the option to set the DRS automation level for individual objects.</p>
VIEW DRS SUMMARY	Select a cluster in the list, then click this link to display a page containing metrics for DRS performance and cluster balance in the selected data center.
SET DRS AUTOMATION	Select a cluster in the list, then click this link to set the level of the DRS automation for the cluster. Note that clusters must be fully automated in order for workload optimization alerts to run actions set in the policies.

History Tab

Displays a graphical depiction of executed manual and automated optimizations for clusters in the selected data center or custom data center, based on parameters you provide.

Table 4-250. Details for History

Option	Description
Selected WLP process drop-down	The optimization action whose details you want to display.
Time duration drop-down	Last <i>n</i> hours - select the time parameter: last 6, 12, 24 hours or last 7 days.
Quick filter	choose a cluster name to search on.
Squares graphic	toggle between viewing processes in icon or circle form.
Circle	Toggle between viewing processes presented in a circle or on a straight line.
Back arrow - reset action.	Reset action.

If you point your cursor to a specific cluster as displayed on the screen, the details of the cluster appear in a tool tip. Click the note card icon on the lower right of the tool tip to go to the Details screen for the cluster. When displayed in the circle format, rings in the circle indicate how much CPU and how much memory was used at any given time. For example, if memory usage was higher than recommended based on your policy settings, the memory circle appears red.

Note the timeline across the bottom of the screen. When you choose parameters, for example, WLP process name, time parameter and cluster name, indicators appear along the timeline, showing when processes were initiated.

To zero in on a specific event, choose a process from the drop-down menu. You can also click points on the marker floating above the timeline, which causes a descriptive tool tip to appear, then double-click the 'Double-click to zoom' icon on the lower right.

If the event you choose includes an actual movement of VMs, you see a blue ball containing the number of VMs moved and showing the direction of the move and starting and ending clusters.

Optimization Potential Tab

When you run Workload Optimization, vRealize Operations Manager runs compatibility checks and excludes those VMs which have constraints, and only optimizes resources of those VMs which can be moved. If you want to see the total potential of your workload optimization, assuming that all VMs can be moved, click the **CALCULATE OPTIMIZATION POTENTIAL** button in the Optimization Potential tab. Optimization Potential disregards the underlying constraints and recommends moves before the compatibility checks. You can download the report to see more details.

If you want to see what can be realistically optimized, click **OPTIMIZE NOW** in the **Operation Status** tab. After you click **OPTIMIZE NOW**, you can download a report to review incompatibilities.

The optimization potential report helps you understand the difference between the optimization achievable when you run **OPTIMIZE NOW** and the total optimization potential.

See also [Example: Run Workload Optimization](#)

Rightsizing

Use this screen to alter the number of CPUs and amount of memory in oversized and undersized virtual machines.

Where You Find Rightsizing

From the Home screen, select **Rightsizing** under Optimize Capacity in the left pane.

Note Click on a data center graphic to display the details for the data center.

How Rightsizing Works

The Capacity Optimization, Reclaim, and Rightsizing features are tightly integrated functions that enable you to assess workload status and resource usage in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed. With this function, you can change CPU size and memory values for oversized and undersized virtual machines to achieve optimum system performance.

When you open the page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To identify possible oversized and undersized VMs in a data center, click its graphic. The area following refreshes to display details about the selected data center.

"Oversized VMs" displays the number of VMs determined to be oversized based on policies previously set. A chart details suggested reductions in the overall number of CPUs and GBs of memory and shows the percentage of total resources the reductions represent. Similarly, "Undersized VMs" indicates the number of VMs considered to be undersized, with a chart listing suggested increases in CPU and memory.

The table at the bottom of the page provides important information about the VMs. Table headings are Oversized VMs and Undersized VMs. VMs under each heading are grouped by cluster. Click the chevron to the left of a cluster name to list all the oversized or undersized VMs, respectively, in that cluster. You can check the box next to one or more VM names and click the **EXCLUDE VM(S)** button to prevent those VMs from being included in a resizing action. You can also select individual VMs to resize before clicking the **RESIZE VM(S)** button.

Run a Rightsize Action on Oversized VMs

Run the action as follows:

- 1 In the table headings, **Select** Oversized VMs.

- 2 **Select** the boxes next to VMs you want to exclude from the action, if any.
- 3 Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.
- 4 **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.
- 5 Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested reductions for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.
- 6 **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

Run a Rightsize Action on Undersized VMs

Run the action as follows:

- 1 In the table headings, **Select** Undersized VMs.
- 2 **Select** the boxes next to VMs you want to exclude from the action, if any.
- 3 Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.
- 4 **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.
- 5 Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested increases for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.
- 6 **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

Table 4-251. Rightsize Options

Option	Description
Select a data center.	Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.
ALL DATACENTERS X	Toggle: click ALL DATACENTERS on the upper right when you want to switch the view to a filtered list of all data centers. Click X to return to a carousel view of data centers.
View:	Filter results to include data centers, custom data centers, or both. Option appears when you select ALL DATACENTERS on the upper right.
Group BY:	Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select ALL DATACENTERS on the upper right.

Table 4-251. Rightsize Options (continued)

Option	Description
Sort by:	Options (Options appear when you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Alarm clock graphic - list data centers/custom data centers by time remaining. ■ Dollar sign - list data centers/custom data centers by potential cost savings. ■ Scales graphic - list data centers/custom data centers by level of optimization.
Select data center or ADD NEW CUSTOM DATACENTER.	Options (Options appear when you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. ■ Select ADD NEW CUSTOM DATACENTER to display a dialog box that enables you to define a custom data center.
Oversized VMs display	Displays the number of VMs identified as oversized, with suggested reductions for vCPU and memory size.
Undersized VMs display	Displays the number of VMs identified as undersized, with suggested increases for vCPU and memory size.
Table of Oversized and Undersized VMs	<p>Tabular representation of the Oversized and Undersized VMs in the selected data center.</p> <p>Click one of the headings - Oversized VMs or Undersized VMs - to refresh the table with data for that heading. The table lists the relevant VMs. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.</p> <p>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.</p> <p>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.</p> <p>Exclude VM(s): the selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.</p> <p>For Oversized VMs:</p> <ul style="list-style-type: none"> ■ RESIZE VM(s): the system displays a dialog box with suggestions for reducing vCPUs and memory. Click the edit icons to change resource size. <p>For Underseized VMs::</p> <ul style="list-style-type: none"> ■ RESIZE VM(s): the system displays a dialog box with suggestions for increasing vCPUs and memory. Click the edit icons to change resource size. <p>SHOW/HIDE EXCLUDED VMS: toggle displays or hides the list of VMs you previously excluded.</p> <p>INCLUDE VM(s): include the selected VMs in the actionable list.</p>

Manage Optimization Schedules

Enables you to set up a regular schedule for optimizing a selected container.

Where You Find Manage Optimization Schedules

At the Workload Optimization screen, select **SCHEDULE** from the pane: Optimization Recommendation

Option	Description
Schedule Name	Meaningful name for the schedule
Time Zone	Choose the time zone for the action
Recurrence	Indicate how often you want the optimize action to run. Complex schedules can be defined, for example, select the Monthly option and choose to run the action on Tuesdays and every other Thursday, beginning on the fifth of the month.
Start on:	Day to start the optimization schedule.
Start at:	Time to start the optimization schedule.
Expire after:	Designate a set number of scheduled runs.
Expire on:	Designate an exact date for the actions to end.

See also [Example: Schedule a Repeating Optimization Action](#)

Workload Automation Policy Settings

Provides options for refining policy settings specifically for Workload Optimization.

Where You Find Workload Automation Settings

Access this screen through the Policies pages:

Select **Administration** from the menu, then select **Policies** from the left pane.

Click **Policy Library**, then click either the **Add New Policy** icon or the **Edit Selected Policy** icon. In the Add or Edit Monitoring policy workspace, on the left click **Workload Automation**.

Refer to [Workload Automation Details](#) .

View DRS Summary

The View DRS Summary page provides insight and perspective into the actions DRS is taking to balance a cluster. You can view DRS settings for the cluster and cluster balance metrics, and determine if recent vMotions are DRS- or user-initiated.

Where You Find the View DRS Summary Page

From the Home screen, select **Workload Optimization** under Optimize Performance in the left pane. Then select a cluster name in the Current Workloads pane. The dimmed View DRS Summary and Set DRS Automation links turn live. Click the link to display the DRS summary information.

Table 4-252. DRS Summary Values

Pane/fields	Value
<cluster name>	Name of the selected cluster
Automation Level	Enabled/Disabled. DRS is running or not.
Migration Threshold	Aggressive/Default/Moderate
Active Memory Used	False/ <i>nn</i> %
Cluster Balance	Shows the variations in the DRS cluster balance metric over time as DRS runs. The graph shows how DRS reacts to and clears any cluster imbalance each time it runs.
Cluster Imbalance	The range of potential imbalance values, as expressed in vCenter DRS metrics.
Total Imbalance	The level of imbalance in a cluster, as measured by vCenter DRS metrics.
Tolerable Threshold	The upper limit of what is tolerable in cluster imbalance. Designated by a green dotted line, this is a vCenter DRS metric.
VM Happiness	A bar graph summarizing the total happy and unhappy VMs in the cluster. For individual VMs, there is a presentation of performance metrics related to its happiness, such as %CPU ready time and memory swapped.
Happy VMs	Total of happy VMs are shown in green. Click in the green zone to show a list of these VMs in the Happy/Unhappy VMs pane to the right.
Unhappy VMs	Total of unhappy VMs is shown in red. To show a list of these VMs in the Happy/Unhappy VMs pane to the right, click in the red zone .
Happy/Unhappy VMs	Lists by name all the VMs in the zone you clicked in the VM Happiness pane.
VM Metrics	Shows the trend in VM happiness or unhappiness
Recent vMotions	The number of recent vMotions, plotted against time.
vMotion Details	Shows the number of DRS-initiated and user (non-DRS) initiated vMotions over time. You can choose which type you want to view.
Date/VM	Date of a given vMotion.
Source/Destination	Source and destination of moved VMs.
Type	DRS-initiated or user initiated.

Optimization Schedules

Use the Optimization Schedules page to edit or delete optimization schedules that you set up in the Manage Optimization Schedule Dialog Box at the Workload Optimization main screen.

Where You Find Optimization Schedules

- From the Home screen, select **Administration > Configuration > Optimization Schedules**.
- At the [Workload Optimization Page](#) page, select in the data center whose optimization schedule you want to edit or delete. Then click **SCHEDULE** in the Optimization Recommendation pane.

Table 4-253. Optimize Schedules Options

Option	Description
Edit icon	Select a schedule from the list, then click the Edit icon. The Manage Optimization Schedules appears, with the data for the selected schedule filled in.
Delete icon	Select a schedule from the list, then click the Delete icon. The selected schedule is deleted and does not run.

See also [Example: Run Workload Optimization](#)

Optimize Placement

A two-page dialog box that provides information about optimizing the workload of a selected container. When you run the optimization action, vRealize Operations Manager checks which of the VMs can be moved to a different cluster for better optimization of resources, based on the settings you entered in your Operational and Business Intent settings. You can download a report that provides information about the list of VMs that were included in, and excluded from, the move plan. The report provides reasons as to why some VMs were excluded from the plan.

First page: The current workload ("before," for example, CPU 105%) and projected results ("after," for example storage utilization 45%) for a possible optimizing action.

Second page: The exact moves planned for compute and storage resources.

Note It is possible that there is no optimization move plan. Review the report to see why vRealize Operations Manager could not provide a move plan.

Where You Find Optimize Placement

At the Workload Optimization screen, select OPTIMIZE NOW in the Optimization Recommendation pane.

Table 4-254. Optimize Clusters Options

Option	Description
Compare Cluster Balance	If you are satisfied with the before and after numbers (First page, above), click NEXT.
Review Optimization Moves	<p>If you are satisfied with the moves planned (Second page, above), click BEGIN ACTION.</p> <p>Note Review the optimization plan report before you click BEGIN ACTION.</p>
Download Report	<p>The optimization plan report is in CSV format, and provides the following information:</p> <ul style="list-style-type: none"> ■ Summary of the optimization plan. ■ Summary of the moves that make up the optimization plan. ■ Issues related to the data center. Resolve these issues before proceeding with the optimization. ■ Issues and incompatibilities applicable to specific VMs and their configurations. Resolve these issues, if applicable. ■ Failed move attempts applicable to the specific VMs and their target destinations, as determined from the VM move plan. Resolve these issues and incompatibilities.

See also [Example: Run Workload Optimization](#).

Predefined Dashboards

5

vRealize Operations Manager includes a broad set of simple to use, but customizable dashboards to get you started with monitoring your VMware environment. The predefined dashboards address several key questions including how you can troubleshoot your VMs, the workload distribution of your hosts, clusters, and datastores, the capacity of your data center, and information about the VMs. You can also view log details.

Each set of dashboards is complemented with a series of out-of-the-box customizable alerts and reports to assist with your operational awareness. Alerts, reports, and dashboards, each have a purpose with minimal overlap. Several activities that are carried out using alerts should be carried out using dashboards. Reports should be kept to a minimum as they are not interactive and do not provide timely information.

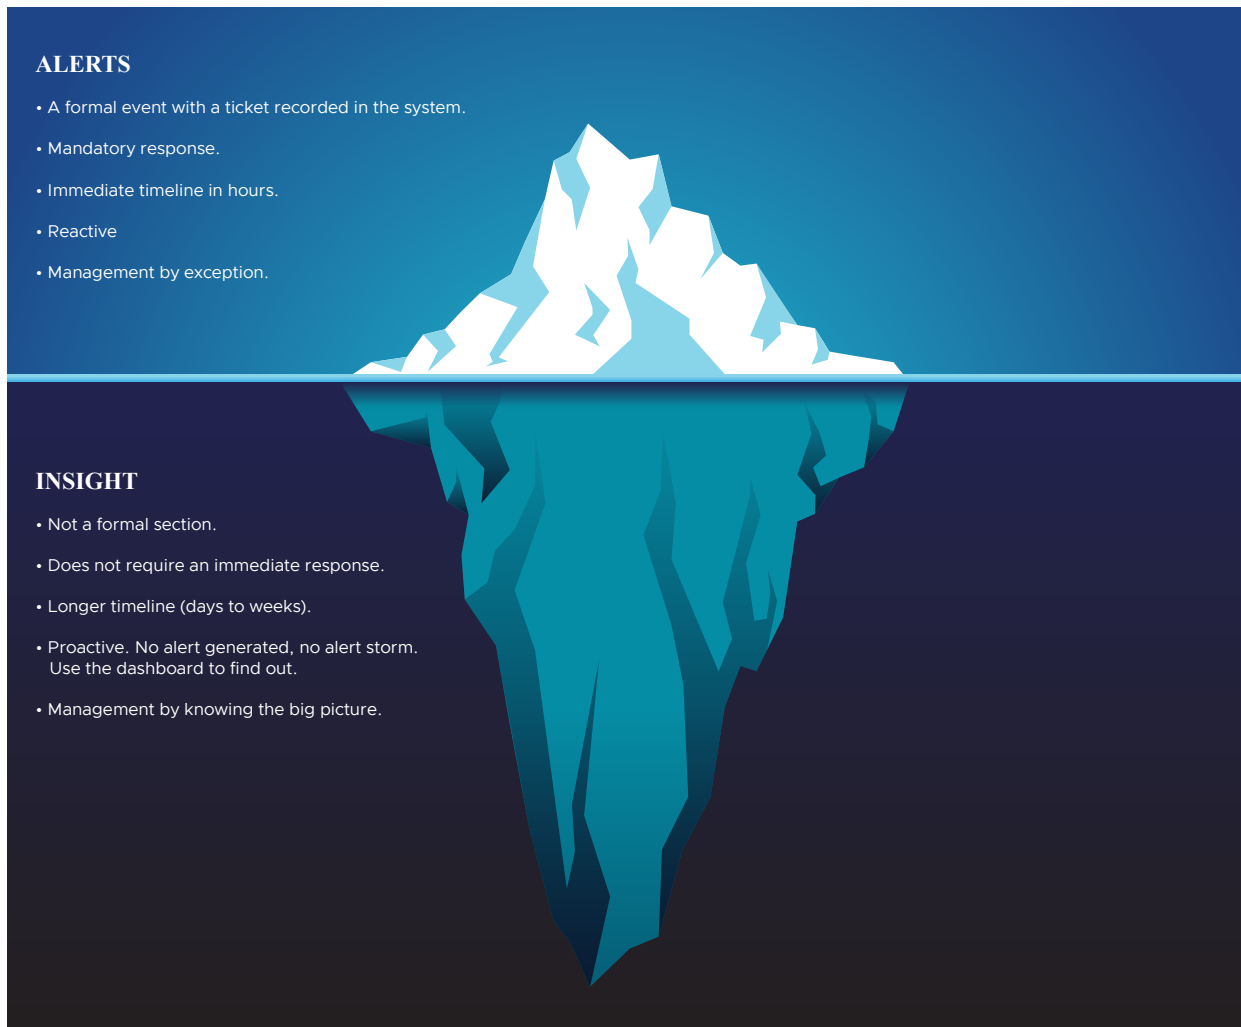
The following table details how alerts, dashboards, and reports are complimentary.

	Alerts	Dashboards	Reports
Nature	Reactive	Proactive	Passive. For those with no access to vRealize Operations Manager/vRealize Operations Cloud and vRealize Log Insight.
Suitability	Exception (something went wrong)	Exception Big Picture Details Analysis	Big Picture Exception (but not urgent) No analysis as it is not interactive
Use Case	Troubleshooting (the start)	Monitoring Troubleshooting (the actual)	FYI (optional) Export for further analysis (spreadsheet)
Time & Urgency	Urgent (minutes) & Important	Regular (daily, SOP)	Not urgent (monthly) & optional No point in a daily report. For daily reports, login for interactivity
Access Requirement	Online. Desktop. 1280 * 1024 pixel	Online. Desktop. 1280 * 1024 pixel	Offline or mobile. Small resolution. Email. Laptop or tablet.
Scope/Area	Availability Performance Compliance Configuration (?) Capacity (less relevant, unless it is an emergency)	Availability Performance Capacity Compliance Configuration Inventory	Same as dashboards, but: <ul style="list-style-type: none"> • without interactivity • time bound (e.g. calendar month) • No performance report, covered in Capacity
Roles	Operations Team	Operations Team Architect Team	IT Management (not hands-on) Auditor (compliance)

Insight vs Alerts

vRealize Operations Manager dashboards support a concept we call insight. Insight complements alerts but does not replace it. Alerts miss the larger picture and only see what is triggered. For one object that reaches the threshold, there might be many just beneath the threshold. The objects below the threshold are called insight.

Alerts might auto-close if the symptoms disappear. Managing alerts is not the same as minimizing alerts. Minimizing alerts is about preventing alerts.



Working with Predefined Dashboards

The default dashboard that appears when you click **Dashboards** in the menu is the **Getting Started** dashboard. You can close a dashboard from the left pane by selecting the dashboard and clicking the **X** icon. The dashboard you last opened is displayed the next time you navigate to **Dashboards** in the menu. If there is only one dashboard left in the left pane, you cannot close it.

To access the predefined dashboards, from the left pane, click the **Dashboards** drop-down menu.

To access the deprecated dashboards, from the left pane, click the **Dashboards** drop-down menu and then select **Dashboard Library > Deprecated**.

You can customize dashboards and widgets if you have vRealize Operations Advanced edition or higher. Any customization you make is overwritten during upgrade and as a result, it is recommended that you back up your dashboards before an upgrade.

This chapter includes the following topics:

- [The Getting Started Page](#)
- [Availability Dashboards](#)

- [Capacity Dashboards](#)
- [Configuration Dashboards](#)
- [Cost Dashboards](#)
- [Performance Dashboards](#)
- [Dashboard Library](#)
- [Software Defined Wide Area Network Dashboard](#)
- [vRealize Automation 8.x Dashboards](#)
- [Service Discovery Dashboards](#)
- [Inventory Dashboards](#)
- [Microsoft Azure Dashboards](#)
- [AWS Dashboards](#)
- [Dashboards in VMware Cloud on AWS](#)
- [Dashboards in NSX-T Management Pack](#)

The Getting Started Page

Operations Management is a set of interdependent disciplines. Knowing the relationship between these disciplines is as important as knowing each of them separately. The relationship between the disciplines matters because the symptom displayed and the root cause are often two different things, for example, sometimes a configuration problem can lead to a performance problem.

Availability

- Availability considers HA (high availability) settings. As a result, planned downtime (for example, ESXi on maintenance mode) impacts availability.
- Availability, done right, does not impact Capacity and Performance as it is already accounted for.
- The higher the Availability SLA, the higher the price. There is a significant difference for each additional 9 of availability. Five 9s costs a lot more than four 9s.

Performance and Capacity

- Performance is more time sensitive and important than capacity. You must manage performance first and then manage capacity.
- Performance and capacity have an opposite relationship. Highest performance is achieved at lowest capacity, as that is when the VM or the infrastructure is delivering the most amount of work.
- Capacity management is about maximizing utilization, without compromising any performance. It also considers latent workload and future demand.

Cost and Price

- Cost goes hand in hand with capacity. The higher the utilization of the IaaS, the lower the cost per VM. Cost is separate from capacity as it can be optimized without reducing capacity.
- Price can move independent of cost. It has concepts such as discount and progressive pricing. Use price to discourage large unused VMs.
- The better the performance SLA, the higher the price the customer is willing to pay, hence the term Price/Performance.

Compliance and Security

- Compliance is measured against both internal and industry standards.
- Security is related, but not the same as configuration.

Configuration and Inventory

- Inventory is related, but not identical to configuration. Configuration impacts performance, cost, capacity, and compliance. Therefore, it is the primary focus of optimization assessment. Inventory is what you have. Configuration includes properties of what you have. For example, the number of VMs in a cluster are a part of the inventory and not a part of configuration. The number of ESXi hosts in a cluster are a part of inventory and configuration because that is how the cluster is designed. The cluster is configured with eight ESXi hosts for the same reason.

There are two types of counters that impact performance and capacity. Contention is the primary counter for performance, and utilization is the primary counter for capacity. Utilization serves performance and capacity differently. For performance, look at the actual and real utilization. For capacity, it is measured against usable capacity (after HA and a buffer). While they have a negative correlation, contention can develop at low utilization. Unbalanced and configurations are two typical causes of low utilization. Allocation complements demand as newly provisioned VMs tend to be idle (which can last for months). Future load cannot be detected by the demand model as they do not exist. The allocation model should be used to complement the demand model.

The Seven Pillars of Operations Management and the Management Process

The best practice of operations management requires you to distinguish between the pillar and process. The pillar is what you must manage, and the process is how you manage them.

What to Manage	Management Dashboards The Seven Pillars of Operations Management.	Day 0: Planning Set your target threshold according to your expectations.	Day 2: Monitoring Compare the reality with the plan.	Day 2: Troubleshooting Identify possible issues and resolve them.	Day 2: Optimizing Reduce cost, increase efficiency, and automate process.
	Availability	Yes	Yes	Yes	May be
	Performance	Yes	Yes	Yes	Yes
	Compliance	Yes	Yes	No	Yes
	Capacity	Yes	Yes	No	Yes
	Cost	Yes	Yes	No	Yes
	Configuration	Yes	Yes	No	Yes
	Inventory	No	May be	No	No

Each pillar is an individual unit of management, namely capacity management, performance management, and compliance management. They represent individual disciplines and are compatible with one another. Each pillar's complexity depends on the technology, for example, vSAN's capacity is more dynamic than the central array. In vSAN, changing the storage policy can create a sudden spike.

Day 0 provides the expected result. Some companies conduct a stress test, load test, so they know what to expect when the real load comes in. Without proper planning, you cannot know what the reality is, as you have not defined the process well.

Troubleshooting is an activity and not something you manage. It focuses on the reason, and then formulates a solution to prevent future incidents. Incidents either mean something dead, slow, or breached. You troubleshoot availability, performance, and security.

Inventory is something you have, not something you plan. You plan for capacity, with a certain configuration. Inventory merely accounts for what you have. Nothing to troubleshoot nor optimize.

Using the Getting Started Page

The Getting Started page breaks tasks into broad three broad categories, Management, Flows, and Collections. Use the Getting Started Dashboard to understand the relationship between these categories.

The Management category includes the seven pillars of operations, Availability, Performance, Compliance, Capacity, Cost, Configuration, and Inventory.

The Flows category of dashboards covers the process that includes Troubleshooting, and Optimization. You can use the Troubleshooting dashboards to resolve any potential issues related to availability, contention, utilization, and configuration. Troubleshooting is more than simply identifying the problem. It focuses on the reason behind the problem and also formulates a solution to prevent reoccurrence. An incident means that something is either dead, slow, or has been breached. You can troubleshoot availability, performance, and capacity. Use the Optimization dashboards to enhance the performance of your environment. You can choose to

correct a problem area, update, simplify, or improve your virtual machines and infrastructure. You can optimize performance, capacity, cost, and configuration. You even improve the availability of your system to an extent but you cannot enhance the compliance or inventory.

Lower Cost	<ul style="list-style-type: none"> • Reclamation: Orphaned VMs, powered off VMs, idle VMs, and oversized VMs snapshots. • Reduce DC Footprint: Save software (MS, Red Hat, VMW), hardware (server, storage, network), and data center (rack, space, cooling, UPS). • Move burst capacity from own to on-demand.
Better Performance	<ul style="list-style-type: none"> • Performance Profiling: Enable proactive monitoring via actual baseline. • Establish performance SLA that complements availability SLA. • NOC Dashboards: Insights followed by alerts. • Faster business service using self-service and approval workflows.
Lower Complexity	<ul style="list-style-type: none"> • Standardize architecture. • Standard operating procedure. • Reduced human error due to automation. • Upgrade outdated software and replace ageing hardware.
Higher Customer Satisfaction	<ul style="list-style-type: none"> • Internal IT department: Reputed among Apps team. • External SP: Repeat business. • Price/Performance: Ability to justify or defend pricing.
Higher Compliance	<ul style="list-style-type: none"> • Internal compliant (for example, vSphere Hardening). • Industry regulation (for example, PCI DSS, HIPAA).

The Collection category comprises of Public Cloud and the Library sections. The AWS and Azure dashboards are displayed under the Public Cloud dashboards. You can choose to view the overall performance of these services or view specific dashboards related to the services. The Library contains dashboards related to the Network Operating Center and the Executive. It also lists dashboards that do not fit into the pillars of operation, like the VOA and the deprecated dashboards.

Using each of these categories you can drill down to the specific use cases and problems you are trying to solve. Each problem statement is associated with a predefined dashboard that you can access through this page. To view a dashboard, click the dashboard type and then select a dashboard from the Getting Started page or click the dashboard name listed on the right side of the Getting Started page.

Note Deprecated dashboards are no longer part of the Getting Started page. They can be accessed from the dashboards drop-down menu under Dashboard Library.

Availability Dashboards

Availability covers the uptime of the object now and the uptime trend over time. The availability of hybrid clouds should be tracked at both the provider and consumer layers to understand the availability of the environment. These dashboards show the current uptime and the uptime percentage over the past month.

VM Availability Dashboard

Use the **VM Availability** dashboard to calculate the availability of the Guest OS. The availability of the Guest OS is calculated because the Guest OS might not be running even when the VM is powered on. There are two layers of Availability, that is, the Consumer layer and the Provider layer. This dashboard covers the Consumer layer. You can view VMs in the selected data center, uptime trend for a selected cluster, and so on.

Design Considerations

The **VM Availability** dashboard helps you check the availability (uptime in percentage) of VMs, as availability is typically part of the services provided by the IaaS provider.

This dashboard does not check the application uptime because it is possible that the application such as, a database, or a web server, is down while the underlying Windows or Linux is up. Generally, the service provided by the IaaS team is only for Windows or Linux. For information on the application, use the network ping or application-specific agent such as application monitoring.

How to Use the Dashboard

- In the **Datacenters** widget, click any data center from the list.
 - To view the overall information, click the **vSphere World** object.
 - The other widgets are automatically updated once you click any datacenter.

- Create a filter that reflects your class of service for this widget. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads and you can focus on the important VMs. You can achieve this by creating a vRealize Operations Manager custom group for each class of service.
- The **VMs by Uptime in the last 30 days** widget displays the average uptime of VMs grouped by their availability. The bucket distribution helps you cater to a wide array of environments. If you are monitoring only production VMs where the uptime is expected to be near 100% all the time, edit the bucket to meet your operational needs.
- The VMs in the **Selected Datacenter** widget display all the VMs that are currently deployed to the data center. The average uptime is displayed for the last month. For a production VM, expect this number to be 100% or closer to 100%.

Note The Services column will be blank unless Service Discovery is enabled and the services/processes are discovered on a specific virtual machine.

- The VMs column includes all VMs including the powered off VMs.
- Click any VM in the **VMs by Uptime in the last 30 days** widget to view the details in the **VM in the Selected VM Powered On Status**, **Selected VM Uptime Trend**, and **Selected Cluster Uptime Trend** widgets.
- The **Selected VM Uptime Trend** widget displays the selected VM's Guest Tool Uptime (%) across the last 30 days.
- The **Guest OS: Services** widget displays the service state over time and the process or services running inside the Guest OS. If Guest OS services or processes are discovered inside a VM, their availability is analyzed. This requires the Service Discovery.
- The **ESXi Host(s) where the VM has run** widget displays the historical migration of the VM. This can be useful in determining the cause of a VM downtime.

Points to Note

- The metric only tracks the availability of VMware Tools and not the entire Guest OS. If VMware Tools is not up, it assumes the Guest OS to be down. You can check that this is not a false negative by adding a few line charts that display the evidence of activity. A good counter is IO counters such as Disk IOPS, Disk Throughput, and Network Transmit Throughput, because IO requires CPU processing. CPU usage is not a reliable counter as the work by VMkernel on the VM is charged to the CPU counters.
- vRealize Operations Manager exhibits a new ping adapter. This allows you to enhance the accuracy of the uptime measurement by creating a super metric that adds the ping information or by checking the process using an agent, such as application monitoring.
- Add a property widget that lists the selected VM properties to give you more context about the VM. In a large environment, the VM name alone might not provide enough context.

vSphere Availability Dashboard

There are two layers of Availability, that is, the Consumer layer and the Provider layer. The **vSphere Availability** dashboard covers the Provider layer. This dashboard includes a cluster and not an ESXi host because the cluster is operationally a single compute provider. This dashboard considers the N+1 design, where the cluster can withstand one host failure. Logically, a cluster with fewer hosts has a higher risk.

Design Considerations

The **vSphere Availability** dashboard helps you analyze and report the uptime, as availability is typically part of the official business SLA. It is also often required in the monthly operational summary report.

This dashboard is not designed for live monitoring of the uptime. A NOC style of dashboard is better suited for those use cases. VMware Tools such as vRealize Log Insight must be leveraged as the fault is typically preceded with soft errors.

How to Use the Dashboard

- The **Clusters** widget lists all the clusters in the environment. It is sorted by the lowest uptime so that the cluster with the lowest uptime in the last one month is displayed.
 - The **Running Hosts** column is color-coded as logically a smaller cluster has a higher risk. A single host failure results in a relatively higher capacity degradation.
 - The **vSAN?** column is hyper-converged, which means both the compute and the storage part is considered.
 - The **Admission Control Policy** column is based on the Cluster Configuration \ DAS Configuration \ Active property. The mapping between the code to name is:
 - -1 : Disabled
 - 0 : Cluster Resource Percentage
 - 1 : Slot Policy (Powered-on VMs)
 - 2 : Dedicated Failover Hosts
 - In a large environment, creating a filter for the list of clusters can make it more manageable. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, you can easily view your gold clusters.
- Click any cluster from the **Clusters** widget.
 - The cluster uptime is automatically plotted in the **Selected Cluster Uptime Trend** widget. It uses 99%, 99.%, and 99.99% as the threshold for red, orange, and yellow colors respectively.
 - The ESXi host details in **ESXi in the Selected Cluster** widget are automatically updated. For more context, you can add a property widget that lists the selected ESXi host properties.

- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.
- The **Datastores not available** widget lists only the datastores with powered off status. This covers both local and shared datastores. To add context, consider adding an extra column such as the data center where it resides, and the datastore types such as NFS and VMFS.
- The **Port Group Availability** widget lists port groups that currently have an uptime of less than 100%. To add context, consider adding an extra column such as the data center where it resides, number of used ports, and the maximum number of ports.
- For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.
- In a large environment, you can create a filter for this dashboard. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads.
- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.

Points to Note

- You can add vCenter Server and NSX components availability. This requires the VMware SDDC Health Monitoring Solution.

Ping Overview Dashboard

Use the Ping Overview dashboard to configure the ping functionality and verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN. You can view ping adapter details like, latency distribution and packet loss distribution in this dashboard.

Customizations Available for Your Use

For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.

Widget Information

- Latency distribution - You can use this widget to see the objects that are experiencing high latency.
- Packet Loss Distribution – You can use this widget to see the objects that are experiencing high packet loss.

- Ping Targets – You can use this widget to view the list of ping targets grouped by their FQDN. Latency and packet loss information is also displayed for the ping objects.
- Breakdown by Source Initiator – You can use this widget to view the List of ping statistics by the source (ping initiator). You can ping the target from multiple locations, to determine if the issue is network-related or server-related.

Capacity Dashboards

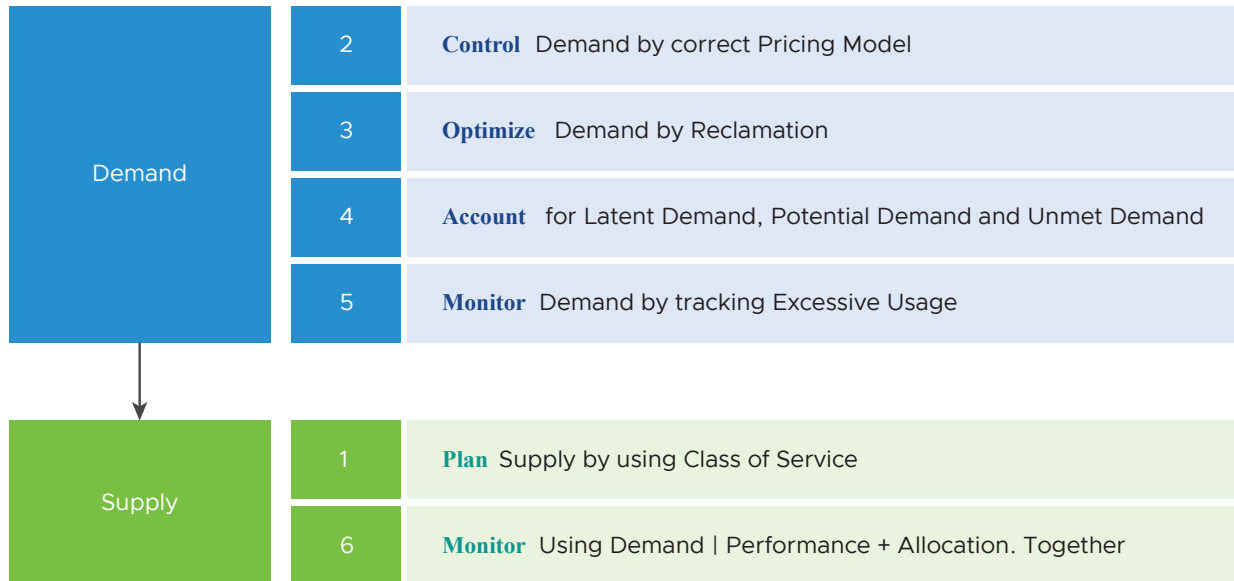
Capacity quantifies the resources used, resources remaining, and opportunities to reclaim unused resources. Projections of the demand provide a proactive view of capacity. The **Capacity Dashboards** display capacity in terms of time remaining before capacity is projected to run out, the amount of capacity remaining, the number of VMs that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

Capacity management is about balancing demand and supply. It is about meeting demand with the lowest possible cost.

For IaaS or DaaS, capacity management begins before the hardware is deployed. It begins with a business plan that defines what class of service will be provided. Each class of service, for example, gold, silver, bronze is differentiated by the quality of service and covers the availability, for example, 99.99% uptime for Gold, 99.95% uptime for Silver. It also covers performance, for example, 10 ms disk latency for Gold, 20 ms disk latency for Silver, and security or compliance.

The quality incurs cost and in turn drives price. Gold VM is higher per vCPU and per GB of RAM because it has a higher quality of service. A proper pricing model must be planned. If you want your customers to rightsize in advance, then a 64 vCPU VM has to be more than 64x the price of 1 vCPU VM. If the pricing model is a simple straight line, there is no incentive to go small and no penalty if it is over provisioned. In this case, you end up forcing rightsizing in production, which is a costly and time consuming process.

Demand is more than the active load that is consuming your capacity. Since capacity based on utilization is incomplete by itself, the principles displayed in the following figure are considered.



- Latent demand. Many critical VMs are protected with Disaster Recovery. During a Disaster Recovery drill or an actual disaster, this load is consumed.
- Potential demand. Many newly provisioned VMs take time to reach their expected demand. It takes time for the database to reach the full size, the user base to reach the target, and the functionalities to complete. When this is achieved, it results in the increase in demand.
- Unmet demand occurs when the VM or Kubernetes Pod is undersized. The load is running nearly 100% most of the time.
- Excessive demand can wreak havoc in a shared environment. A group of highly demanding VM can collectively impact overall performance of the cluster or datastore.

Cluster Capacity Dashboard

The **Cluster Capacity** dashboard helps you visualize information differently by providing choices for customization. Use this dashboard to highlight the clusters that need attention. The **Cluster Capacity** dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology.

Design Considerations

See [Capacity Dashboards](#) for common design considerations among all the dashboards for capacity management. The **Cluster Capacity** dashboard considers the following factors that impact capacity:

- Contention
- Utilization

- Allocation
- Reclamation

Contention is included as it directly measures the performance. If your cluster is unable to serve its existing workload, then do not add a new workload. By definition, if the cluster does not have room for a new workload, then its capacity is full. The ideal scenario is that the cluster must run at 100% utilization but 0% contention. In this case, the cluster is productive and your investment is well used.

Utilization is the primary counter for capacity, as it reflects the actual live usage of the resources. When utilization is high, it does not matter if the overcommit ratio is far below your target because the cluster is full. Also, the utilization must not be very low.

Allocation complements utilization as not all workload is real. Some demands can suddenly emerge such as:

- Newly provisioned VM
- Disaster Recovery
- Undersized VM
- Auto-scale VM (a group of web servers behind a Load Balancer)

Reclamation is included as it can impact your decision and the wastage can be common. Capacity can be low, but if you can reclaim a sizeable chunk of wastage, you can defer the purchase of the hardware.

Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **Cluster Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

- The **Clusters by Capacity Remaining** and **Clusters by Time Remaining (days)** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.
- The two bar charts work together. The ideal situation is low capacity remaining and high time remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

- The three heat maps are **Time Remaining**, **Capacity Remaining**, and **VM Remaining**.
- The cluster size is made constant for ease of use. If your cluster sizes are not standardized, consider using the number of ESXi hosts to display the difference in sizes.

The third layer displays a table, accompanied by other widgets to display details of the selected cluster.

- **Clusters Capacity List** widget. If any cluster needs attention, then select the cluster to view the related details.
- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.
- Reservation can impact the efficiency of your cluster. If your cluster size varies, complement the reservation number by displaying a relative value.
- Number of VMs is displayed because the newly provisioned VMs might not be active yet. They are often mistaken as idle, as they can remain unused for months. When you view VM increasing but the demand remaining low, it is a sign of potential demand coming up in the future.
- Workload can be low, but is the overcommit ratio high? The newly provisioned VMs tend to be idle for weeks, and suddenly increase. Use the **VM Count** widget to view if there was recent growth.
- You can check why it is low on capacity. Is it because of real workload or just reservation?

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.
- If you have screen real estate, add a cluster size information. Add cluster size. Small clusters are less efficient from a capacity perspective due to higher overheads and the inability to support larger VMs.
- The peak is defined as the highest among any ESXi hosts. If the peak is higher than the cluster-wide average, then it is unbalanced and is a common reason for suboptimal capacity. You can add a peak to complement the average utilization. Find out the cause of unbalance and optimize it.
- Add peak to complement average utilization. This lets you focus on unbalance, a common reason for suboptimal capacity. Find out the source of unbalance, which can be an opportunity for optimization.
- This dashboard is not designed for the stretched cluster as it requires its own capacity model.

Datastore Capacity Dashboard

The **Datastore Capacity** dashboard highlights the datastores that need attention. This dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology. The **Datastore Capacity** dashboard is designed for both the VMware administrator and the Storage administrator to foster closer collaboration between the two teams.

Design Considerations

See [Capacity Dashboards](#) for common design consideration among all the dashboards for capacity management.

Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

- The **Shared Datastores by Capacity Remaining** and **Shared Datastores by Time Remaining** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.
- The two bar charts work together. The ideal situation is low capacity remaining and high time remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

- There are three heat maps, the primary heat map being **Remaining Capacity**.
- The other two heat maps **Used Capacity by Datacenter** and **Used Capacity by Datastore Cluster**, cover used capacity. The **Used Capacity by Datastore Cluster** heat map is designed for datastore clusters.

The **Shared Datastores** widget is grouped by data center.

- If you use Datastore Cluster as your standard, replace the grouping with it. This widget is sorted by the least capacity remaining.

Select a datastore from the **Shared Datastores** widget. The remaining widgets will automatically display the capacity details of the selected datastore.

- The **Disk Space** widget displays the total capacity allocated and the actual capacity used. You can compare the total Capacity vs the Provisioned capacity vs the Used capacity. If the allocated space increases and the actual capacity does not, then it means that the VMs have not been used. You can ensure that you do not run out of space sooner than expected.
- In the **VM Count** widget, a rising number that is not complemented by a similar rise in the used space indicates a latent demand.
- There are three reclamation opportunities: powered off VM, snapshot, and orphaned VMDK.
 - The snapshot must be 0 GB. If it is not 0, then it should be temporary. A snapshot lasting beyond a day must be investigated.
 - Orphaned VMDK are the ones that are not associated with any VM. The orphaned VMDK must be 0.

Points to Note

Storage in VMware IaaS is presented as a datastore. In a large environment, group datastores as datastore clusters for ease of operations. vSAN uses datastores to present its storage, but it requires a different formula for capacity and performance management. In certain situations, RDM (Raw Device Mapping) and network file shares are also used by certain VMs.

ESXi Capacity Dashboard

The **ESXi Capacity** dashboard supports the **Cluster Capacity** dashboard and is also required for the non-clustered ESXi.

Design Considerations

See [Capacity Dashboards](#) for common design considerations among all the dashboards for capacity management.

How to Use the Dashboard

The **Summary** heat map provides an overall view of the ESXi Host capacity, grouped by their clusters.

- Each ESXi host is represented by a box, displaying their capacity remaining.
- The ESXi host size is made constant for ease of use. If your ESXi sizes are not standardized, consider using the number of physical cores or Total CPU GHz to display the difference in sizes. Ensure that the smallest ESXi is not too small.
- Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

The **ESXi Hosts Capacity** widget lists all the ESXi hosts in your environment, grouped by their parent cluster.

- The standalone ESXi are displayed at the bottom under No Group.
- In a large environment with many data centers, you can zoom into a specific vCenter or data center. You can also filter or search for specific ESXi hosts matching certain names.
- The **99th Percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric. To rule out the outlier, the worst performance (which is equivalent to 100th percentile) is not considered. Also, the performance threshold is set to be stringent.

Select one of the ESXi hosts from the **ESXi Hosts Capacity** widget. All the three line charts automatically display the trend of selected ESXi host.

- Displays both total and usable utilization in terms of RAM and CPU.
- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.
- A technology refresh is often used to address the capacity shortage. Consider adding a property widget that displays the hardware model and specification to help you determine the age of the hardware.

VM Capacity Dashboard

The **VM Capacity** dashboard provides a quick overview of all the VMs in a given data center and their capacity and time remaining. This dashboard is designed for the Capacity team and not for the Operations team. It provides a long term and a top-down view, enabling the Capacity team to plan the future expansion and refresh of the aging hardware technology.

Design Considerations

See [Capacity Dashboards](#) for common design considerations among all the dashboards for capacity management.

The **VM Capacity** dashboard helps you visualize information differently by providing choices for customization. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

Select a data center from the **Datacenters** widget.

- The **VMs by Capacity Remaining** bar chart displays the distribution of VMs by capacity remaining in the selected data center. It provides a quick overview of the VMs that are undersized or oversized.
- The **VMs by Capacity Remaining** heat map provides details by grouping the VMs by clusters, so that you can view the clusters that need attention.
- The VM size has been standardized for better visualization. You can add the size that suits your capacity team better.
- The **VMs Capacity in the Selected DataCenter** widget is sorted by the VM with the least capacity remaining. You can sort it by Time Remaining to suit your capacity team better. This table is color-coded.

Select a VM from the **VMs Capacity in the Selected DataCenter** widget. All the remaining widgets automatically display the capacity information of the selected VM.

- The **Disk** widget displays capacity at the Guest OS partition level. There is no overall capacity at VM level because different partitions have different capacity.

Points to Note

- Use the custom property and add more context to the VM, such as, owner name, clusters where the VM is running, and datastores where the VM files are stored.

VM Reclamation Dashboard

The **VM Reclamation** dashboard helps you manage various types of reclamation that can be done on virtual machines. This dashboard is designed for the Capacity and the Operations team. The reclamation is grouped by buckets. Use this dashboard to view the trend charts that help you analyze the growth over time without changing the context.

Design Considerations

The **VM Reclamation** dashboard helps you visualize information differently by providing choices for customization. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

This dashboard is divided into two sections:


- The first section provides a summary, giving you an overall picture of reclamation.
- The second section provides details, giving you the actual VM name to reclaim.

Review the summary information.

- The summary is presented in the three bar charts, **Count of VM by Snapshot Size**, **Count of Powered Off VM by Disk Space**, and **Count of Idle VM by Memory Footprint**. Each bar chart corresponds to an area you can reclaim.
 - The **VM Snapshots** widget corresponds to VMs that are more than a few days old.
 - The **Powered Off VM** widget assumes they are back up and it is safe to delete them.
 - The **Idle VM** widget helps you reclaim memory but not CPU. The idle VM memory occupies ESXi host's physical memory.
- Idle VM does not display any CPU usage as there is nothing to reclaim and no CPU is being used. Since the CPU is idle, the only benefit is the overcommit ratio.
- Memory reclamation is based on the memory footprint at the parent ESXi host. The value inside the Guest is not what is being reclaimed, and so it is irrelevant.
- Adjust the bucket size to suit your operational requirements.
- Select any of the VMs in the above widgets to view its trend over time. The trend chart is placed on the same page, so you can review it without opening a new screen. This helps you quickly toggle between VMs.
- If the snapshot is expanding rapidly, ensure that the VM disk is large (relative to the underlying datastore) as it can fill up the datastore.

Points to Note

There are five areas of reclamation, so start with the easiest first.



- Non-VM indicates files are not associated with VM. Typically, these are ISO files.
- Orphaned file is a file in the datastore that is no longer associated with any VM. For orphaned Raw Device Mapping (RDM), look from the storage array if there is any ESXi host mounting it. The orphaned VMDK is not listed in this dashboard as it is not associated with a VM. If your environment has orphaned VMDK, add a fourth column in this dashboard.
- Snapshots are not back up and they cause performance problem to the VM. Keep them only for protection during change. Once the change is validated as successful, keeping the snapshot does a disservice to the VM.
- If your environment is large, change the dashboard filter to a functional filter. Group by the class of services such as gold, silver, and bronze and default the selection to the least critical environment. In this way, you can be active in reclamation.
- If reclaiming is a long drawn manual process in your organization, add a filter by department or VM owners. One way to do this is to create a vRealize Operations Manager custom group.
- If the VM name in your environment does not provide sufficient business context, add more information in the table to give context to the VM. Information such as VM owner, clusters where the VM is running, and datastores where the VM files are stored can be useful in the analysis.
- Disk cannot be reclaimed immediately. They have to be in the powered-off stage at least for a week.

vSAN Capacity Dashboard

The **vSAN Capacity** dashboard complements the vSphere **Cluster Capacity** dashboard by displaying capacity related to vSAN. To manage vSAN capacity, use both dashboards.

Design Considerations

As this dashboard is designed to complement the vSphere **Cluster Capacity** dashboard, it shares the same design consideration. It focuses on the storage and vSAN specific metrics but does not list non-vSAN clusters.

See [Capacity Dashboards](#) for common design considerations among all the dashboards for capacity management.

Wastage is displayed by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

How to Use the Dashboard

The **vSAN Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

The first layer displays two distribution charts.

- The **Clusters by Capacity Remaining** and **Clusters by Time Remaining (days)** bar charts summarize the clusters based on capacity remaining and time remaining. Just because you are running low on capacity does not mean you are running out of time.
- The two bar charts work together. The ideal situation is low Capacity Remaining and high Time Remaining. This means that your resources are cost effective and are working as expected.

The second layer displays a heat map.

- The three heat maps are **Time Remaining**, **Capacity Remaining**, and **VM Remaining**.
- The cluster size is made constant for ease of use. If your cluster sizes are not standardized, consider using the number of ESXi hosts to display the difference in sizes.

The third layer displays a table, accompanied by other widgets to display details of the selected cluster.

- **vSAN Clusters** widget. If any cluster needs attention, then select the cluster to view the related details.

Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.

vSAN Stretched Clusters

The vSAN Stretched Clusters dashboard provides an overview of the cluster resources used across vSAN fault domains. Using the stretched clusters dashboard you can monitor the resource consumption at the site level for Preferred Sites and Secondary Sites. You can create custom dashboards for specific vSAN stretched cluster metrics.

Where to View vSAN Stretched Cluster Objects

On the menu, click **Dashboard > Capacity and Utilization > vSAN Stretched Clusters**.

You can also view the vSAN stretched cluster objects from **Environment > VMware vSAN > vSAN and Storage Devices > vSAN Clusters**, if the vSAN cluster is a stretched cluster.

The vSAN Stretched Clusters dashboard provides information about CPU Capacity, Cores, Memory Capacity, and Disk Capacity for the Preferred Site and the Secondary Site. You can identify the vSAN stretched clusters running out of capacity looking at the utilization metrics.


Configuration Dashboards

As an operations management software, vRealize Operations Manager focuses on the impact to day-to-day operations a product has, rather than the feature of the product itself. Products under monitoring, such as vSphere and vSAN, can have features that are related, but have a different impact on operations. For example, vSphere provides Limits, Reservation, and Shares for the VM.

Limits, Shares, and Reservation. As a feature, they are closely related, and appear in the same dialog box and must be learned as one. However, they impact operations differently. The following table describes that in more detail.

VM Limit	Impacts the VM	<ul style="list-style-type: none"> • Should not be used. Right-size instead. • Results in unpredictable performance of the Guest OS. 	Incorrect Configuration
VM Reservation	Impacts the Infrastructure	<ul style="list-style-type: none"> • Keep the total amount low, and relative to the total capacity of the cluster. • Absolute value. A 2-GHz reservation is in fact a 2-GHz reservation. • Results in suboptimal infrastructure capacity, as overcommit is not possible. 	Suboptimal Configuration
VM Share	Impacts the Infrastructure	<ul style="list-style-type: none"> • Keep the number of variations to below three. One for each class of service. • Relative value. 2000 worth of reservations depends on the value of other VM reservations. Be careful when you move the VM to another cluster, as the relative value changes. • Results in complex operations. It is harder to troubleshoot performance when the dynamic entitlements of each VM fluctates more. 	Complex Configuration

vRealize Operations Manager follows the principle that there are different impacts on operations and applies a methodology for looking at configuration. It does not group the settings by features or objects. Rather, it begins with the impact and prioritizes what can be done.



Each operation is unique and as a result, customers run operations differently. What is right for other customers, might not be right for you. Even in the same environment, what is right for a development environment might not be appropriate for a production environment.

The following table lists some of the areas for improvement for the operations in your environment:

Areas of Improvement

	Correct it?	Update it?	Simplify it?	Optimize it?
IaaS Consumer: <ul style="list-style-type: none"> Process Applications Guest OS Container VM 	<ul style="list-style-type: none"> Java JVM or Database \ memory config too large relative to Guest OS Guest \ Metric not collecting Guest \ High TX Broadcast packets VM \ Tools not installed VM \ Tools not running VM \ CPU Limit VM \ Memory Limit VM \ Old Snapshot VM \ On local Datastore 	<ul style="list-style-type: none"> Guest OS \ Tools Guest OS \ Windows Guest OS \ Linux Guest OS \ Telegraf agent VM \ Hardware (vmx) 	<ul style="list-style-type: none"> VM \ Large VM (CPU, RAM, Disk) VM \ lots of disks, NIC card VM \ lots of IP address. VM \ with RDM VM \ on multiple datastores VM \ Fault Tolerant VM \ SRM protected VM \ Hot Add/Remove \ CPU VM \ Hot Add/Remove \ RAM 	<ul style="list-style-type: none"> Java JVM or Database \ memory config too small relative to Guest OS Guest OS \ no visibility Container \ smaller than the parent VM VM \ Tools unmanaged VM \ bigger than the whole ESXi cores. VM \ bigger than CPU socket. VM \ Large Snapshot VM \ Reservation.
IaaS Provider: <ul style="list-style-type: none"> Telegraf ESXi Cluster Datastore & Cluster Switch and Port Group Hardware NSX vSAN 	<ul style="list-style-type: none"> ESXi \ vMotion disabled ESXi \ Disconnected from vCenter ESXi \ Maintenance Mode ESXi \ NTP disabled ESXi \ Standalone Cluster \ Admission Control disabled Cluster \ HA disabled Cluster \ HA Failover % Cluster \ DRS disabled Cluster \ DRS manual Cluster Inconsistency <ul style="list-style-type: none"> BIOS, ESXi: version BIOS, ESXi: Power Management ESXi Storage Path ESXi Hardware Datastore Cluster inconsistency <ul style="list-style-type: none"> Capacity Performance Datastore \ single path Datastore \ no path. This is unlikely. NSX \ no redundancy for Controller, Manager 	<ul style="list-style-type: none"> ARC \ server ARC \ agent ESXi \ hardware ESXi \ vSphere ESXi \ 1 Gb NIC. Server \ not on warranty vCenter \ version Datastore \ VMFS version vSAN \ version Switch \ version NSX \ version 	<ul style="list-style-type: none"> ESXi \ Too many variations. No standard Cluster \ Many VM Shares (CPU) Cluster \ Many VM Shares (RAM) Cluster \ Resource Pools Cluster \ Stretched compute + storage Cluster \ 32 nodes or more Cluster \ VM to Host affinity Cluster \ Too many storage paths Datastore \ Shared by >1 cluster WLP uses this Datastore \ Many paths Network \ LBT? Network \ MAC Address change 	<ul style="list-style-type: none"> ESXi \ low CPU cores count ESXi \ low RAM size ESXi \ Powered Off ESXi \ HT Disabled ESXi \ 4 socket or higher. Cluster \ small clusters \ host especially for vSAN Cluster \ small clusters \ CPU Cluster \ small clusters \ RAM Cluster \ EVC Mode Cluster \ High Reservation Cluster \ DRS Automation Level Cluster \ DPM disabled vSAN \ All Flash: Dedupe disabled vSAN \ All Flash: Compressed disabled Datastore \ small Datastore \ low VM count Datastore \ no ESXi Distributed Switch \ unused

Design Considerations

The dashboards display configurations that need immediate attention, before displaying the overall configuration. This helps you take measures toward optimizing configuration.

Operations vary among customers, and as a result, it is not possible to design one dashboard to meet every customer's operational needs. A configuration that is important for one customer might not be relevant for another customer. Tailor the dashboard to your unique environment. You can collapse or expand the widgets to allow relevant data to be displayed.

The overall layout is designed to balance ease of use, performance (loading time of the dashboard page), and completeness of configuration check. As a result, not all configuration settings are displayed. Lack of screen real estate is another consideration behind the design.

Cluster Configuration Dashboard

Use the **Cluster Configuration** dashboard to view the overall configuration of vSphere clusters in your environment, especially configurations that need attention.

Design Considerations

See the [Configuration Dashboards](#) page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

The **Cluster Configuration** dashboard is organized into sections for ease of use.

- The first section of the dashboard consists of three bar-charts. They correspond to the three main features of vSphere clusters, namely High Availability (HA), Dynamic Resource Scheduler (DRS), and Distributed Power Management (DPM).
 - HA: The best practice is to enable HA admission control. You can specify the admission control policy in the vCenter Server and the threshold for failover shares.
 - DRS: The best practice is to enable DRS. Envision the vSphere cluster as a single logical computer that balances within itself.
 - DPM: The best practice is to enable DPM in an environment where environmental concern is the top priority or the high peak rarely occurs as most of the time you run very low utilization.
- The second section of the dashboard consists of eight pie-charts. They show the relative distribution of key configurations.
 - Two of the bar-charts cover admission control. You must enable admission control. The pie-charts displays the policy code instead of the policy name, as it is based on the property: `Cluster Configuration | Das Configuration | Active Admission Control Policy`. The mapping between the code to name is:
 - -1 = Disabled
 - 0 = Cluster Resource Percentage
 - 1 = Slot Policy (Powered-on VMs)

- 2 = Dedicated Failover Hosts
- There are two bar-charts that cover the HA Failover Share. One for CPU, and one for memory.
- The next two bar-charts cover DRS settings. You might want to fully automate DRS, which means that there is no operator intervention required for both initial VM placement and subsequent load balancing, but with a moderate migration threshold (value = 3.0). The value ranges from 1.0 to 5.0.
- There are two pie-charts that show reservation. One for CPU and one for memory. Minimize the total reservation value as it prevents overcommit of resources and hence results in less optimal utilization. Memory reservation can remain and occupy the memory space of the ESXi host, even though the VM does not use the memory anymore. Consider the analogy of unused files that you have not opened for months in the c:\ drive of your laptop. They still take up space on the hard disk. Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service.
- The third section of the dashboard consists of two bar-charts. They show the absolute distribution of the clusters.
 - The first bar-chart displays the cluster grouped by the number of ESXi hosts. Small clusters, defined as having a lower number of ESXi hosts, have a higher overhead while large clusters have a higher risk if there are cluster-wide outages. Performance risk is lower, because there are more nodes that DRS can tap on, but if there is an actual problem, troubleshooting can be tougher, because there are more nodes to analyze. For large clusters, have a disaster recovery plan as an unexpected cluster-wide outage can impact many VMs.
- The fourth section of the dashboard lets you drill-down to an individual cluster.
 - A table lists all the clusters with their key configuration. You can export this list as a spreadsheet for further analysis or reporting.
 - Select a cluster. The list of ESXi hosts under the cluster, with shares and resource pool information, is automatically filled up.
 - Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service. Avoid providing different services to individual VMs as that increases the complexity of the cluster performance.
 - Keep the number of resource pools minimal.
 - Some of the columns are color coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

ESXi Configuration Dashboard

Use the **ESXi Configuration** dashboard to view the overall configuration of the ESXi hosts in your environment, especially the configurations that need attention.

Design Considerations

See the [Configuration Dashboards](#) page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

The **ESXi Configuration** dashboard is organized into sections for ease of use.

- The upper section of the dashboard displays basic ESXi configurations that should be standardized for ease of operations.
 - There are six pie-charts that are displayed as one set because there is a relationship between their values. There should be a correlation between them. Ideally, the ESXi version, the ESXi build, and the BIOS must be identical across all ESXi hosts in a cluster. Keep the variations of the hardware model, NIC speed, and storage path minimal. The more complex the pie-chart, the more variants you have. This results in complex operations, that potentially results in higher operating expenses.
 - The configurations should reflect your current architecture standard. Each pie-chart counts the occurrences of a particular value. A large slice signifies that the value is the most common value, and if that is not your current standard, then you must address it.
- The second section of the dashboard displays configurations that are potentially suboptimal.
 - The three bar-charts display various size dimensions of the ESXi hosts. The bar-charts are designed to be seen as one set. Ensure that there are a minimal number of variations to reduce complexity.

- Smaller ESXi hosts have a relatively higher overhead, and are limited in running larger VMs. If they have a low core count, they might be using an outdated CPU. Small ESXi hosts are more expensive on a per core, per GB, per rack unit basis than larger ones if they occupy the same space. However, a 4-CPU socket ESXi host is likely to be too large, resulting in a concentration risk (too many VMs in a single ESXi host). Maintain a good balance that balances your budget and risk constraints.
- Adjust the distribution chart bucket size to fit your environment.
- The third section of the dashboard displays configurations that you might want to avoid.
 - The six bar-charts focus on security, availability, and capacity settings that you can set as a standard. For example, you should consider enabling the NTP daemon for a consistent time, which is critical for logging and troubleshooting.
 - The three tables list the actual ESXi hosts that are in a non-productive state. They can be on maintenance mode, powered off, or in a disconnected state.
- The last section of the dashboard displays all the ESXi hosts in your environment.
 - You can sort the columns and export the results into a spreadsheet for further analysis.
 - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result, there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- For complete visibility, consider adding physical server monitoring by using the appropriate management pack. For more information, see the following [page](#).

Network Configuration Dashboard

Use the **Network Configuration** dashboard to view the overall configuration of vSphere distributed switches in your environment, especially for the areas that need your attention.

Design Considerations

See the [Configuration Dashboards](#) page for common design considerations among all the dashboards for configuration management.

How to Use the Dashboard

The dashboard is organized into two sections for ease of use.

- The first section displays network configurations that need your attention.
 - There are five bar-charts that focus on critical security settings.
 - The last bar-chart displays the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section provides overall configuration information, with the ability to drill down to a specific switch.
 - Click the row to select a switch from the list.
 - The ESXi hosts, port groups, and the VMs on the switch are displayed.
 - Review each of the tables. For the ESXi host table, ensure that the settings are consistent.
 - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
 - You can sort the columns and export the result into a spreadsheet for further analysis.

Points to Note

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- For complete visibility, consider adding physical network device monitoring by using the appropriate management pack. For more information, see the following [page](#).
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

VM Configuration Dashboard

Use the **VM Configuration** dashboard to view the overall configuration of virtual machines in your environment, especially for the areas that need attention.

Design Considerations

See the [Configuration Dashboards](#) page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

How to Use the Dashboard

- Click the row to select a data center from the data center table.
 - In a large environment, loading thousands of VMs increases the web page loading time. As a result, the VM is grouped by data center. In addition, it might make sense to review the VM configuration per data center.
 - For a small environment, vSphere World is provided, so you can view all the VMs in the environment.

The **VM Configuration** dashboard is organized into three sections for ease of use. All the three sections display the VM configuration for the selected data center.

- The first section covers limits, shares, and reservations.
 - Their values can easily become inconsistent among VMs, especially in an environment with multiple vCenter Servers.
 - Shares should be mapped to a service level, to provide a larger proportion of shared resources to those VMs who pay more. This means that you should only have as many shares as your service levels. If your IaaS provides gold, silver, and bronze, then you should have only three types of shares.
 - Value of the shares and reservation is relative. If you move a VM from one cluster to another (in the same or different vCenter Server), you might have to adjust the shares.
 - Reservation impacts your capacity. Memory reservation works differently from CPU reservation, and it is more permanent.
- The second section covers VMware Tools.
 - VMware Tools is a key component of any VM, and should be kept running and up to date.
- The third section covers other key VM configurations.
 - Keep the configurations consistent by minimizing the variants. This helps to reduce complexity.
 - **VM Network Cards** widget. If you suspect that your environment might have a VM with no NIC, consider adding it as a dedicated bucket.
- The last section of the dashboard is collapsed by default.
 - You can view all the VMs with their key configurations.
 - You can sort the columns and export the results into a spreadsheet for further analysis.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- `No data to display` does not imply that there is something wrong with data collection by vRealize Operations Manager. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.
- The pie-chart and bar-chart cannot drive other widgets. For example, you cannot select one of the pie-slices or buckets, and expect it to act as a filter to a list or a table.
- You can apply a specific color in a pie-chart or distribution chart for a specific numeric value, but not string value. For example, you cannot apply the color red to the value `Not Installed`.

vSAN Configuration Dashboard

The **vSAN Configuration** dashboard provides overall configuration details and is useful in large clusters with many vSANs, where you have to follow a certain standard configuration.

Design Considerations

See [Configuration Dashboards](#) for common design considerations among all the dashboards for configuration management.

How to Use the Dashboard

The **vSAN Configuration** dashboard is organized into three sections for ease of use.

- The first section displays six pie-charts.
 - There are five bar-charts that focus on critical security settings.
 - The last bar-chart shows the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section displays three bar-charts.
 - The three bar-charts together provide a good overview of the vSAN key capacity configuration. By analyzing the distribution, you can identify if you have capacity configuration that is outside your expectation.
- The last section of the dashboard displays all the vSAN clusters with their key configuration.
 - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
 - You can sort the columns and export the result into a spreadsheet for further analysis.

Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.

- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

Workload Management Configuration Dashboard

This dashboard provides a quick configuration summary of all the key objects associated with workload management such as Supervisor Clusters, Namespaces, vSphere Pods and Tanzu Kubernetes clusters. It is essential that the configuration is consistent across all the objects. Configuration drifts may result in inconsistent performance or availability of the applications leveraging workload management Kubernetes constructs.

Use the dashboard to ensure that the configuration is consistent across all objects.

You can view the following widgets in the dashboard.

- **Environment Summary**
- **Supervisor Cluster Versions**
- **Cluster Status**
- **Pod Data**
- **Supervisor Cluster Configuration Summary**
- **Pod Configuration Summary**
- **Kubernetes cluster Configuration Summary**
- **Namespace Configuration Summary**

Consumer \ Correct it? Dashboard

The **Consumer \ Correct it?** dashboard complements the main VM configuration dashboards by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The dashboard is designed to focus on VMs that need attention. Lists are used to keep it simple, and show actual objects. The lists can be tailored using the filter and the custom group. The lists can also be exported for an offline discussion.

The dashboard is extendable, reflecting the reality that different customers have a different set of settings to verify. Since the dashboard layout is a collection of tables (List View), you can extend it by adding more tables. You can add more List View widgets to verify the VM configurations that your operations require.

How to Use the Dashboard

The **Consumer\Correct it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Tools Widgets:
 - Using VMware Tools has multiple benefits. For the list of benefits, refer to [KB 340](#).
 - vRealize Operations Manager uses VMware Tools to retrieve Guest OS metrics. Without this, right-sizing VM memory can be inaccurate, because the hypervisor metrics (VM Memory Consumed and VM Memory Active) are not designed to measure Windows or Linux memory utilization. ESXi VMkernel does not have visibility into the Guest OS for security reasons.
 - Independent software vendor (ISV) support is the most common reason that VMware Tools is not installed. The ISV vendor might claim that no additional software is installed in their appliance unless they have certified it. For more information about VMware Tools, see the [VMware Tools documentation](#).
 - If VMware Tools is installed, there might be reasons why the application team disables it. The Infrastructure team should inform and educate their application team, and document the technical recommendations about why VMware Tools is recommended to be running all the time.
- CPU Limits and Memory Widgets:
 - It is recommended that you do not use memory and CPU limits as it can result in an unpredictable performance. The Guest OS is not aware of this restriction as it is at the hypervisor level. It is recommended that you shrink the VM instead.
- Guest OS Counters Missing Widget:
 - There is no visibility into the Guest OS performance counters because the requirements are not met. The memory counter is especially important as VM Consumed and VM Active are not replacements for Guest OS counters. See [KB 55675](#) for more details.
- Old Snapshot Widget:
 - Ensure that the snapshot is removed within one day after the change request. If not, it might result in a large snapshot and impact the performance of the VM.

Points to Note

- Add a banner summary to the top of this dashboard so that you can verify if there is an incorrect confirmation. Add a scoreboard and select the World object and then collapse all the tables below. Create a super metric for each summary and apply it to the World object.

- In a large environment, create a filter for this dashboard to enable you to focus on a segment of the environment. Group it by a class of service such as, gold, silver, and bronze. Default the selection to gold, your most important environment. In this way, your monitoring is not cluttered with less critical workloads.
- There are other VM configurations that maybe relevant to your environment. Review the list of VM settings that you might want to add to this dashboard.
- For context, add a property widget that lists the selected VM properties. In this way, you can check the property of your interest without leaving the screen. Multiple List View widgets can drive the same property widget, so you do not have to create one property widget for each List View.
- If your operations require it, add a list of VMs that do not have these three key performance counters: CPU Run Queue, CPU Context Switch, and Disk Queue Length.

Consumer \ Optimize it? Dashboard

The **Consumer \ Optimize it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Optimize it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities. A suboptimal configuration might not impact performance or increase complexity, but it can be more expensive.

Design Considerations

The **Consumer \ Optimize it?** dashboard follows the same design considerations specified for the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How To Use the Dashboard

The **Consumer \ Optimize it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- VM Reservation:
 - VM reservation causes a positive impact on the VM, but a negative impact on the cluster. Total reservation cannot exceed cluster capacity. This creates a suboptimal cluster as VMs do not use the entire assigned memory at the same time.
 - VM reservation places a constraint on the DRS placement and HA calculation. Avoid using reservation as a means to differentiate performance SLA among all the VMs in the same cluster. It is difficult to correlate CPU Ready with CPU Reservation. A VM CPU Ready does not improve two times because you increase its CPU reservation by two times. There is no direct correlation.

- Guest OS visibility:
 - Since your workloads are sharing resources and are over-committed, your operations are easier if you know what is running inside. This helps with monitoring and troubleshooting, resulting in optimal operations.
 - For critical VMs, consider logging the Guest OS, such as Windows and Linux, to capture errors that do not surface as metrics. These errors typically appear as events in the log files or in the event database in the case of Windows. Use vRealize Log Insight to parse Windows events into log entries that can be analyzed.
- Snapshot:
 - Old snapshots tend to be larger. They consume more space and have a higher chance of impacting performance.

Points to Note

See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and as a result shares limitations and customization ideas.

Consumer \ Simplify it?

The **Consumer \ Simplify it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Simplify it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Consumer \ Simplify it?** dashboard follows the same design considerations specified for the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The **Consumer \ Simplify it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Large VMs (CPU, Memory, and Disk):
 - A large VM, relative to the underlying ESXi host and datastore, requires more careful planning (Day 0) and monitoring (Day 2).

- Ensure that the VM size does not exceed the size of the underlying ESXi host. If your ESXi host has CPU hyper-threading, do not count the logical processor. Instead, count the physical core. For best performance, keep it within a (non-uniform memory access) NUMA boundary.
- During monitoring, verify if the VM is highly utilized. If the VM vCPU count is equal to the ESXi cores, and the VM is running at almost full capacity, you might not be able to run other VMs. Large VMs can impact the performance of other VMs, especially if it is given higher shares. Only when the large VM is under-utilized, can the ESXi hosts run other VMs.
- If the number of configured vCPUs on a VM is higher than the number of cores per socket on the ESXi, the VM can experience the NUMA effect. If the ESXi has more than one physical CPU (socket), cross-NUMA access negatively impacts performance.
- The larger the VM, the longer the time required to vMotion, Storage vMotion, and backup.
- For disk space, if the disk is thin-provisioned and under-utilized, you can deploy other VMs in the same datastore. Ensure that the snapshot is tracked closely, as the risk of capacity running out is higher for a large virtual disk.
- VMs with many virtual disks:
 - It is simpler to have a 1:1 mapping between Guest OS partitions and the underlying virtual disk (VMDK or RDM).
 - For performance and capacity, evaluate the disks and partitions. Each virtual disk must be monitored in terms of IOPS, throughput, and latency. Having multiple virtual disks increases the monitoring and troubleshooting need.
 - If the reason for having many virtual disks is performance, identify which counter serves as proof that multiple virtual disks are required. It is possible that the performance required is met by a single virtual disk.
- VM with many IP addresses or NICs:
 - A VM might need multiple networks, such as production, back up, and management. It is recommended that you route the network interfaces through the NSX-Edge VM. A VM that has multiple network interfaces can bridge the network, causing security risks or network problems.
 - A VM that is part of multiple networks can do so with just a single NIC. A single NIC can be configured to access multiple networks, with each interface having their own IP configuration.

Points to Note

See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

Consumer \ Update it? Dashboard

The **Consumer \ Update it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Update it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Consumer \ Update it?** dashboard follows the same design considerations specified for the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The **Consumer \ Update it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated Tools Widget:
 - Lists all the VMware Tools versions that are still supported. Tailor the filter to fit your operational needs.
- Outdated VM Hardware Widget:
 - Lists all the VM vmx versions that are not 13, 14, 15, or 16. Tailor the filter to fit your operational needs.
- Outdated Windows and Red Hat Widgets:
 - Lists all the Windows client versions that are not version 10.
 - Lists all the Windows server versions that are not versions 2016 and 2019.
 - Lists all the RHEL versions that are not version 7 or 8.
 - If you run other operating systems like Ubuntu, clone the widget. You can also repurpose the widget if you do not run RHEL and Windows.

Points to Note

See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

Provider \ Correct it? Dashboard

The **Provider \ Correct it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is

designed for vSphere administrators and the platform team. The **Provider \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

Design Considerations

The **Provider \ Correct it?** dashboard follows the same design considerations specified in the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations.
 - A cluster is the smallest logical building block for compute. Consider it as a single computer with physically independent components. As a result, consistency matters.
 - Clusters with DRS set to manual. This means that DRS initiated vMotion does not take place unless the administrator manually approves it. Since DRS calculates every five minutes, your quick approval is required to prevent a change of condition.
 - Clusters with HA disabled. Without high availability provided by the infrastructure, each application must protect itself from an infrastructure failure.
 - Clusters with DRS disabled. DRS focuses on performance and capacity, while HA focuses on availability. Without DRS, you must build a buffer on every ESXi host to cope with peak demand.
 - Clusters with Admission Control disabled. Reservation is respected only when Admission Control is enabled.
- The second section covers the ESXi host configurations.
 - ESXi with Network Time Protocol disabled. Logs are a critical component of operations, and are the main source of information in troubleshooting. While troubleshooting performance across objects, the sequence of logs determines which event is the likely root cause, as the oldest event starts the chain of events.
 - A disconnected ESXi host indicates that the ESXi host is not participating in HA and you cannot migrate any VM on it.
 - An ESXi host that is in maintenance mode does not contribute resources to the cluster or the data center if there is a standalone ESXi.
- The third section covers ESXi host configurations that must be consistent within a cluster.
 - BIOS version and ESXi versions.
 - BIOS Power Management, ESXi: Power Management. Ideally, should be set to OS controlled. The ESXi level should be set to balance level.
 - ESXi Storage Path. Ensure that the number of paths and the path policies are identical.

- ESXi hardware specifications. Different specifications can result in inconsistent performances experienced by the VM.

Points to Note

- See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- If you have a standalone ESXi, and you plan to replace it with a clustered ESXi host, add a table to list them.
- Based on your security settings, add a table to check the Distributed Switch and Port Group to ensure that security settings such as promiscuous mode, are used correctly.

Provider \ Optimize it? Dashboard

The **Provider \ Optimize it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Optimize it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

Design Considerations

The **Provider \ Optimize it?** dashboard follows the same design considerations specified in the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations:
 - A small cluster has a higher HA overhead when compared to a large one. For example, a three-node cluster has 33% overhead while a 10-node cluster has 10%. For vSAN, a low number of hosts limits the availability option. Your choice of FTT is relatively more limited.
 - Many small clusters result in silos of resources. As a cluster behaves like a single computer, ensure that it has enough CPU cores, CPU GHz, and Memory. For ESXi in 2020, it is typical to have 512 GB of RAM. This results in 12 TB of RAM for a 12-node cluster, which is enough for DRS to place many VMs as it balances them.
 - If there is a lot of reservation, add a list for clusters with a relatively high reservation. If your clusters are of different sizes, use a super metric to convert the reservation value to a percentage.

- The second section covers ESXi host configurations.
 - Small ESXi. A small host faces scalability limits in running a larger VM. While a 2-socket, 32-cores, 128 GB memory ESXi can run 30 vCPU, 100 GB RAM VMs, the VM experiences a non-uniform memory access (NUMA) effect.
 - ESXi powered off. You can mark the ESXi hosts for decommissioning using the custom property feature of vRealize Operations Manager. You can then create a separate list, so they are not overlooked.
- The third section cover storage and network.
-
- Unused network (distributed port group). This is a potential security risk as you might not monitor it.

Points to Note

- See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- For CPU cores, a change in vSphere licensing means that the ideal core is 32-cores per CPU socket. This maximizes the software license. For more information, see the vSphere [Pricing Model](#).

Provider \ Simplify it? Dashboard

The **Provider \ Simplify it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Simplify it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

Design Considerations

The **Provider \ Simplify it?** dashboard follows the same design considerations specified in the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboard

- Click the row in the **Clusters** widget to select one of the clusters from the table.
 - A cluster is more complex to operate when it has resource pools, shares, and limits.
- Review the list of resource pools:
 - Ensure that the number of VMs in each resource pool reflects the intended settings for the VM. The resource pool value is divided and shared among the VMs. The more the VMs, the lesser the resources allotted to each VM.

- Verify if there are VMs who are siblings to the resource pools.
- Verify if the resource pools are further split into subresource pools.
- Review the CPU Share and Memory Shares pie-charts:
 - Multiple combinations of shares, especially both CPU and memory, makes troubleshooting difficult.
 - Each share must map to exactly one class of service, such as one for gold and one for silver as the shares define the class of service. Shares are also relative, meaning the value depends on the value of sibling objects, such as, resource pool or VM. Ensure that the values are consistent across clusters to avoid unintended consequences while moving the VM to another cluster.
- Review the CPU Reservation and Memory Reservation tables:
 - High total reservation, especially both CPU and memory, complicates the cluster operations as it impacts the HA slot calculation, and limits the DRS choice of placement.
- Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

Points to Note

See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

Provider \ Update it? Dashboard

The **Provider \ Update it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Update it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

As part of operations best practices, keep the infrastructure up to date. Running outdated components that are too far behind the latest version, can cause support problems or upgrade problems. It is common that the fix for the problem is only available in the later versions. Outdated hardware can also result in higher operating costs. Outdated hardware might cost more data center footprint, such as rack space, cooling, and UPS. Refreshing your technology and consolidation are two common techniques to optimize cost.

Design Considerations

The **Provider \ Update it?** dashboard follows the same design considerations specified in the [Consumer \ Correct it? Dashboard](#) dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

How to Use the Dashboards

The **Consumer \ Update it?** dashboard is a collection of tables (List View) that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated vSphere Components Widgets:
 - Lists all the vCenter Servers versions that are not 6.7 or 7.0.
 - Lists all the ESXi host versions that are not 6.5, 6.7, or 7.0.
 - Lists all the vSAN ESXi host versions that are not 6.7 or 7.0. A more stringent filter is applied for vSAN because of a relatively higher maturity in the latest release. From vRealize Operations Manager and vRealize Log Insight, there are more counters, properties, and events that improve monitoring and troubleshooting.
 - Lists all the vSphere distributed switches, regardless of the version.
 - You should tailor the filter to fit your operational needs.
- Outdated Server BIOS Widget:
 - Lists all the ESXi hosts regardless of the BIOS version. Edit the widget and tailor the filter to fit your operational needs.
- Other than customizing the existing widgets, consider adding the following checks:
 - ESXi hosts with outdated hardware, using a filter based on your environment.
 - ESXi hosts that are no longer on warranty. Create a custom property to capture the end of warranty.
 - Physical storage arrays with outdated firmware, model, and an expiring warranty.
 - Physical network switch with an outdated OS version and hardware model

Note Install the relevant management pack for the last two points.

Points to Note

See the **Points to Note** section as specified in the [Consumer \ Correct it? Dashboard](#) dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

Cost Dashboards

The dashboards in the cost category cater to cloud administrators who are responsible for managing the expenses related to your cloud infrastructure. Using Cost dashboards, you can compare the cost of VMware cloud infrastructure with other cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently.

Assess Cost Dashboards

The **Assess Cost** dashboard provides an overview of the scale of your infrastructure in terms of physical capacity available.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- You can view the total cost of ownership per month for the infrastructure and the savings opportunities details, if any, for the infrastructure.
- You can view the details of the division of infrastructure investments across all data centers. The dashboard provides the magnitude of each data center in terms of the number of physical servers and virtual machines. It also provides details about the amount of savings that can be achieved from each of these data centers.
- The dashboard displays data about how you invest across clusters of different quality offered across all vCenter Servers.

Base Rate Analysis Dashboard

The **Base Rate Analysis** dashboard helps you analyze the cost efficiency of your data center.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- The total cost of ownership is the cost required to run your data center per month. This is derived from the cost drivers.
- The average cost per VM is derived by considering the cost of all the VMs in your environment. The cost of each VM depends on the base rate of the cluster the VM is placed on and its utilization. The base rate of the cluster is computed based on the total cost of ownership and the expected utilization levels of the cluster. Storage base rates are directly obtained from cost drivers.
- If the cluster is running on an allocation-based capacity model, the base rate is derived from the total cost of the cluster and the over-commit ratio. The base rate is indicative of how costly a resource is, on a given cluster.
- A base rate is derived from the total cost and the expected utilization of the cluster.
- A deeper analysis of the base rates can be performed using the CPU, memory, or storage-related widgets, which help rank clusters and datastores relative to their base rates.

Datacenter Cost Drivers Dashboard

The **Datacenter Cost Drivers** dashboard provides the cost of different data centers in a private cloud.

Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

Widget Information

- You can select individual data centers to view summary and trends. The summary of the data center costs is grouped into two:
 - Compute. Covers all the costs that are spent on compute related hardware, software, and services.
 - Non-Compute. Covers storage and network.
- Expense trends provide cost variations over a period which indicate infrastructure additions or removal to the data center.
- Cluster expenses indicate the component clusters of a data center that consume the costs. Datastores that represent the storage part of the data center cost are listed alongside.

Note Network costs are mapped directly to ESXi hosts and hence are costed under compute as well, as of today. This might change in the future.

- When you select a cluster, you can view the component hosts that the cluster is made up of and their monthly depreciated costs. It also provides details on the purchase cost of the server and how many months until it depreciates completely.

Note Server costs can be suggested out-of-the-box by the system, or can be customized by the user. Depreciation information is not available for servers when the server costs are suggested out-of-the-box by the system. Depreciation information is available for those servers when the server cost is customized by the user.

Showback Dashboard

The **Showback** dashboard helps you navigate between multiple groups that you might want to perform showback for. Select an object to view total costs of the group and potential savings within the group.

Customizations Available for Your Use

The definition of a group is based on several constructs such as the vCenter folder, the vRealize Automation 7 business group, and the vRealize Automation 8 project among others. To change this to your definition, edit the widget and select the desired object types under **Output Filter > Basic > Object Types**.

Widget Information

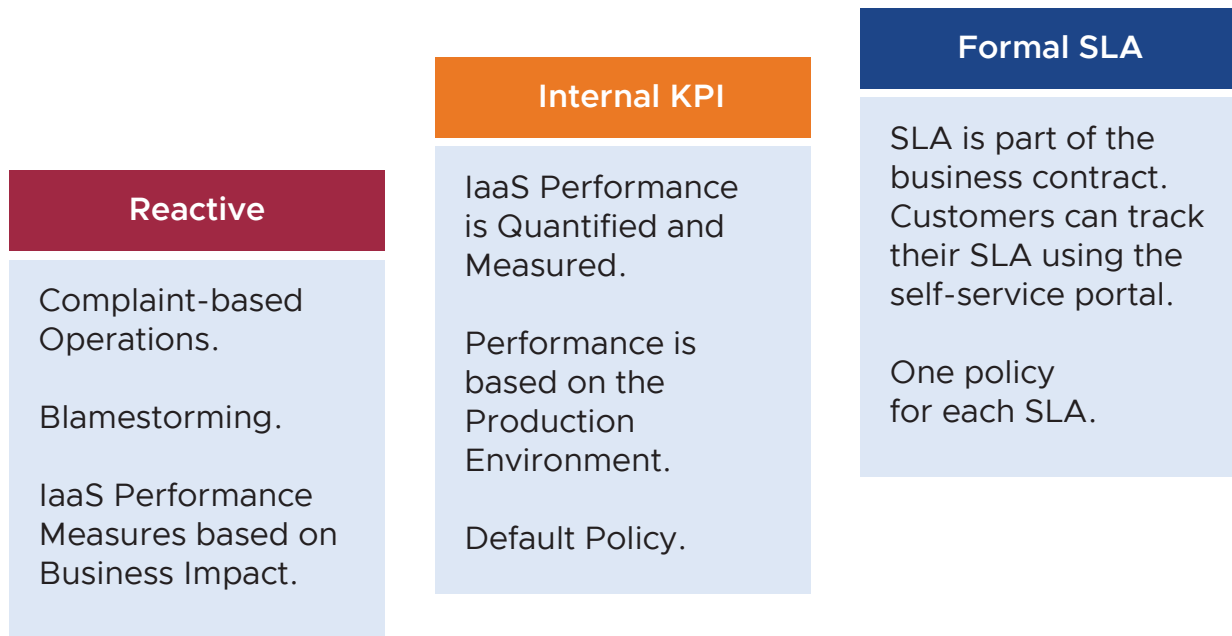
- The **Top Expensive** and **Top Savings** widgets point out the most costly VMs and the VMs with the most savings opportunities within the group.
- A detailed list of the members helps you perform a deeper analysis. You can also export the details into a CSV for offline analysis.
- When you select a particular VM in the list, you can analyze the cost trends for the VM, and as a result, find out the root cause for VM cost variations.

Performance Dashboards

Performance is about ensuring workloads get the necessary resources. Key Performance Indicators (KPI) can be used to identify performance problems related to workloads. Use these KPIs to define SLAs associated with tiers of service. These dashboards use KPIs to display the performance of workloads at the consumer layer and the aggregate performance of workloads at the provider layer.

SLA is the formal business contract that you have with your customers. Typically, SLA is between the IaaS provider (the infrastructure team) and the IaaS customer (the application team or business unit). Formal SLA needs operational transformation, for example, it requires more than technical changes and you might need to look at the contract, price (not cost), process, and people. KPI covers SLA metrics and additional metrics that provide early warning. If you do not have an SLA, then start with Internal KPI. You must understand and profile the actual performance of your IaaS. Use the default settings in vRealize Operations Manager if you do not have your own threshold, as those thresholds have been selected to support proactive operations.

The following graphics depict the above relationship.



The Three Processes of Performance Management

In performance management, there are three distinct processes.

- **Planning.** Set your performance goals. When you architect a vSAN, you must know how many milliseconds of disk latency you want. 10 milliseconds measured at the VM level (not the vSAN level) is a good start.
- **Monitoring.** Compare the plan with the actual. Does the reality match what your architecture was supposed to deliver? If not, you must fix it.
- **Troubleshooting.** When the reality is not according to the plan, you must fix it proactively and not wait for issues and complaints.

To understand what is not healthy for performance management consider the following areas in the given order.

- 1 **Contention:** This is the primary indicator.
- 2 **Configuration:** Check the version incompatibilities.
- 3 **Availability:** Check for soft errors. vMotion stun time, lock up. This requires Log Insight.
- 4 **Utilization:** Check this in the end. If the first three parameters are good, you can skip this.

The Three Layers of Performance Management

There are three main realms of enterprise applications. Each of these realms has its own set of teams. Each team has a set of unique responsibilities and requires the associated skill set. The three realms comprise of Business, Application, and IaaS. Refer to the graphic below to understand the three layers and the typical questions asked on each layer.

Layers		Sample Metrics
Business	Business Result	<ul style="list-style-type: none"> • How many sales did we make today? • How many customers bought our product this week? • On an average, how long did the XYZ transaction take in this hour? • How many customers logged in yesterday? • On an average how long did customers stay logged in?
	Business Transaction	
Application	Individual Node	<ul style="list-style-type: none"> • How long did the SQL query ABCD take in the last 7 days ? • One hour ago, what was the value of the SQL server free memory ? • What is the overall application uptime? • Are my applications configured for performance?
	The System	
IaaS	VM or Container	<ul style="list-style-type: none"> • What is the Windows CPU Run Queue? • In the past 24 hours, what was the peak VM CPU contention? • What was the total number of IO hitting vSAN from 9am - 6pm yesterday? • What is the buffer in a physical switch right now?
	Virtual Infra	
	Physical Infra	

Vertical Metrics depend on each application and its needs

2

Horizontal Common metrics are applicable for all applications

1

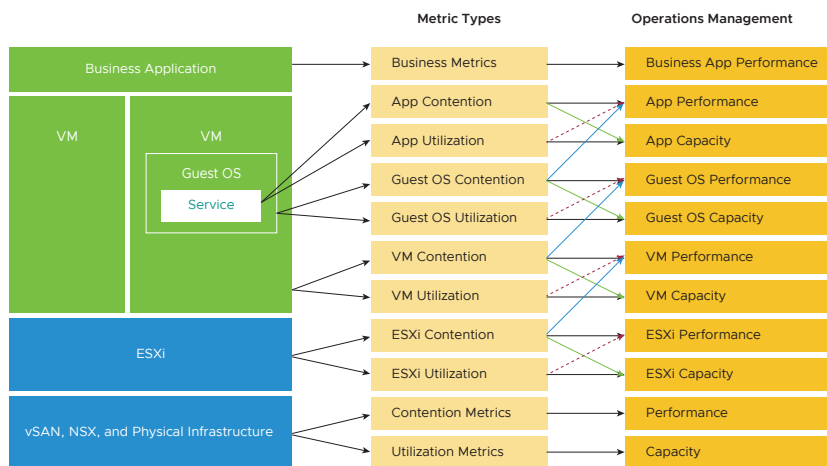
Performance Management is largely an exercise in elimination. The methodology slices each layer and determines if that layer is causing the performance problem. Hence it is imperative to have a single metric to indicate if a particular layer is performing or not. This primary metric is aptly named Key Performance Indicator (KPI).

The upper layer depends on the layer below it, and hence the infrastructure layer is typically the source of contention. As a result, focus on the bottom layer first, as it serves as the foundation for the layer above it. The good part is this layer is typically a horizontal layer, providing a set of generic infrastructure services, regardless of what business applications are running on it.

The Two Metrics of Performance Management

The primary counter for performance is contention. Most look at utilization, because they fear something wrong might happen if utilization is high. That something is contention. Contention manifests in different forms like, queue, latency, dropped, canceled, and context switch.

However, do not confuse ultra-high utilization indicators as a performance problem. If your ESXi host experiences ballooning, compression, and swapping, it does not mean that your VM has a performance problem. You measure the performance of the host by how well it serves its VMs. While performance is related to the ESXi host utilization, the performance metric is not based on the utilization, instead it is based on contention metrics.



It is possible for VMs in the cluster get affected from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at an individual consumer (VM). The following table shows various possible reasons.

Infra Configuration	VM and Guest OS Configuration
<p>ESXi Settings</p> <ul style="list-style-type: none"> ■ Host and BIOS power management causes Frequency to drop. ■ HT enabled. It looks like twice the capacity, but it is actually 1.25 X throughput. ■ ESXi - HW compatibility. Driver and firmware are two areas that can impact the performance. ■ Mismatch of queue depths along the various storage stacks. Must calibrate all the way to the physical array. ■ vMotion too slow or high stunned time. 	<p>VM: Limit, Share, and Reservation</p> <ul style="list-style-type: none"> ■ Make sure that no limit is set. CPU ready includes limit. ■ Make sure that the shares are consistent (as per what the VMs want or you agree to.) ■ Avoid reservation if possible. This impacts the net available resources for the other VMs.
<p>Network</p> <ul style="list-style-type: none"> ■ MTU mismatch. ■ Hops. Especially horse-shoe, or going through multiple ESXi. 	<p>Size: NUMA effect. VM spanning NUMA nodes.</p>
<p>Cluster Settings</p> <ul style="list-style-type: none"> ■ Inconsistent configuration among hosts in a cluster. EVC Mode can play a part if the hosts are from different generations. ■ Resource Pool <ul style="list-style-type: none"> ■ Make sure the shares match the number of VMs. ■ Make sure that no VM is siblings to RP. ■ VM- Host Affinity. ■ DRS Setting. 	<p>Snapshot. IO is processes 2x. VM drivers.</p>
<p>vSAN</p> <ul style="list-style-type: none"> ■ The host where the storage was having performance issues. 	<p>Windows or Linux process ping pong, process runaway, and OS level queue.</p>

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

VM Performance

Since VM is the most important object in vSphere, it warrants an extra explanation. The graphic below lists the counters you should look at.

	CPU	RAM	Network	Disk
Inside Guest OS (Linux, Windows) Need VMware Tools	Run Queue Context Switch	Paging Rate (MB/s) Committed %	OS Output Queue Length Driver Queue	OS Queue Driver Queue
	Utilization	In Use Modified + Standby	Throughput (Mbps) Latency	Latency
Outside Guest OS (Guest OS can't control)	Run I Used System + VMX + MKS	Active, Consumed, Granted, Swapped-in	Throughput	IOPS, Throughput (Large Block)
	Ready + Co-Stop + Overlap IO Wait + Swap Wait	Contention	TX Dropped Packet Normalized Latency	Outstanding IO Latency

The KPI counters can get technical for some users, so vRealize Operations include a starting line to get them started. You can adjust the threshold, once you profile your environment. This profiling is a good exercise, as most customers do not have a baseline. The profiling requires an advanced

	Metric	Green	Yellow	Orange	Red
Guest OS Contention	Total CPU Run Queue	0 - 5	> 5	> 10	> 20
	CPU Context Switch Rate	0 - 5K	< 25K	< 100K	> 100K
	Total Disk Queue Length	0 - 25	> 25	> 50	> 100
Guest OS Usage	RAM Free (MB)	> 512 MB	> 256	> 128	≤ 128
	RAM Page-in Rate (KB/s)	0 - 25K	> 25 K	> 50 K	> 100K
VM Contention	CPU Co-Stop (%)	0 - 2.5%	> 1	> 3	> 5
	[SLA] CPU Ready (%)	0 - 2.5%	> 2.5	> 5	> 7.5
	Total CPU Overlap (ms) at VM level	0 - 1000	> 1000	> 2500	> 5000
	CPU IO Wait	0 - 1000	> 1000	> 2500	> 5000
	[SLA] RAM Contention (%)	0 - 1%	> 1	> 2	> 4
	[SLA] Disk Latency (ms)	0 - 10 ms	> 10	> 20	> 40
	[SLA] Network TX Dropped Packet	0	> 0	> 1	> 2
	CPU Usage (%)	0 - 85%	> 85	> 90	> 95

edition.

Performance Metrics

vRealize Operations Manager uses the following threshold for internal KPI.

IaaS	VM Counter	Threshold
CPU	Ready	2.5%
RAM	Contention	1%
Disk	Latency	10 ms
Network	TX Dropped Packet	0

The table is an example of a stringent threshold. A high standard for performance is used because it is an internal KPI for the consumption of the infrastructure team. It is not an external formal SLA that is confirmed with the customers. There must be a buffer between the internal KPI and the external SLA so that the operations team receive early warnings and has the time to react before the external SLA is breached. A high standard also works from the mission critical point to view to the development environment. If the standard is set to the least performing environment, then it cannot be applied to the more critical development.

A single threshold is used to keep the operations simple. This means that the performance in production is expected to have a higher score than the development environment. The development environment performance is expected to be worse than the production environment, while everything else is equal. A single threshold helps to explain the difference in Quality of Service (QoS) provided by a different class of service. For example, if you pay less, you get a poor performance and if you pay half the price, expect to get half the performance.

The four elements of IaaS (CPU, RAM, Disk, and Network) as mentioned in the table, are evaluated on every collection cycle. The collection time is set at five minutes as it is an appropriate balance for monitoring. If SLA is based on one minute, it is too close and results in either cost increase or reduction in threshold.

Design Considerations

All the performance dashboards share the same design principles. They are intentionally designed to be similar, as it is confusing if each dashboard looks different from one another, considering they have the same objective.

The dashboards are designed with separate two sections: summary and detail.

- The summary section is typically placed at the top of the dashboard to provide the overall picture.
- The detail section is placed below the summary section. It lets you drill down into a specific object. For example, you can get the detailed performance report of any specific VM.

In the detail section, use the quick context switch to check the performance of multiple objects during performance troubleshooting. For example, if you are looking at the VM performance, you can view the VM-specific information and the KPIs without changing screens. You can move from one VM to another and view the details without opening multiple windows.

The dashboard uses progressive disclosure to minimize information overload and ensure the webpage loads fast. Also, if your browser session remains, the interface remembers your last selections.

Many of the performance and capacity dashboards share a similar layout since there is a shared commonality between these pillars of operations.

Guest OS Performance Profiling Dashboard

Use the **Guest OS Performance Profiling** dashboard to know the actual performance of your environment.

Some counters directly impact the performance of Windows or Linux, the operating systems running inside the VM. These KPIs are outside the control of the hypervisor.

Modern operating systems such as Linux and Windows use memory as cache, since it is faster than a disk. Some counters directly impact the performance of Windows or Linux. These KPIs are outside the control of a hypervisor, which means that the ESXi VMkernel cannot control the increase or decrease of the KPI values. The KPI visibility also requires an agent, such as VMware Tools. As a result, they are typically excluded in performance monitoring.

Since they are closer to the applications, it is critical to know their values and establish an acceptable range. The acceptable level of these KPIs among all the VMs in your environment varies. By profiling the actual performance across time and from all VMs, you can establish a threshold that is supported by facts. Since there are 8766 instances of 5 minutes in a month, profiling 1000 VM over a month means you are analyzing 8.8 million datapoints.

Design Considerations

The dashboard uses progressive disclosure to minimize information overload and ensures that the webpage loads fast.

In a large environment, loading thousands of VMs increases the loading time of vRealize Operations Manager. As a result, the VM is grouped by data center. For a small environment, vSphere World is provided so you can see all the VMs in the environment.

How to Use the Dashboard

Select data center from the data centers list. The three tables listing CPU, memory, and disk will show the VMs in the selected data center or vSphere world. Each table shows the highest value in the last one week (2016 datapoints based on five minutes collection cycles), and hence uses the term max as a prefix, for example Max Page-Out/sec or Max Guest OS Disk Queue.

Select any of the VMs in any of the tables. The three line charts are displayed. They are showing data from the same VM to facilitate correlation.

■ CPU table widget:

- The Max CPU Queue column shows the highest number of processes in the queue during the given period. As a best practice, keep the queue below three for each queue. A VM with eight CPUs has eight queues, hence keep this number below 24.
- The CPU Hyperthreading gives twice the queue as it should as both threads are interspersed in the core pipeline.
- CPU Context Switch. There is a cost associated with the context switch. There is no guidance for this number, and it varies widely.

■ Memory list widget:

- In memory paging, the modern operating systems (Linux and Windows) use memory as cache, it is much faster than a disk. It proactively pre-fetches pages and anticipates future needs (Windows calls this Superfetch). The rate pages that are being brought in and out can reveal memory performance abnormalities. A sudden change, or one that has sustained over time, can indicate page faults. Page faults indicate that pages are not readily available and must be brought in. If a page fault occurs too frequently, it can impact application performance. While there is no concrete guidance, as it varies by application, you can view a relative size. operating systems typically use 4 KB or 2 MB page sizes.

■ Disk list widget:

- Disk queues are queued IO commands that are not sent to the VM. They have been retained inside the Guest OS (either at a kernel level or a driver level). A high disk queue in

the guest OS, accompanied by low IOPS at the VM, can indicate that the IO commands are stuck waiting on processing by Windows/Linux. There is no concrete guidance regarding these IO commands threshold as it varies for different applications. You should view this with the Outstanding Disk IO at the VM layer.

Points to Note

- These Guest OS widgets do not appear unless the vSphere pre-requisites are met. For more information, see KB article [55697](#).
- Once you determine an acceptable threshold for your environment, consider adding thresholds to the table so you can easily view the VMs that exceed a threshold.
- The CPU queue is the sum from all virtual CPUs. A larger VM can tolerate a higher queue as it has more processors. If you want to compare VMs of different sizes, create a super metric that calculates the queue per vCPU. For more information, see [Create a Super Metric](#).
- Group the VM by clusters of the same class (for example, Gold), so you can see the profile for each environment.
- For a smaller environment, consider changing the table from listing data centers to listing clusters.

Network Top Talkers Dashboard

Use the **Network Top Talkers** dashboard to monitor network demand in your IaaS. In a shared environment, a few VMs generating excessive activity can impact the entire data center. While a single VM might not cause a serious problem, a few of them can.

Design Considerations

The **Network Top Talkers** dashboard helps you analyze how hard these VMs hit your IaaS. It classifies the workload into two: short bursts and sustained hits. A short burst lasts for a short period, maybe for a few minutes. A sustained hit can last for an hour and cause serious problems.

The **Network Top Talker** dashboard forms a pair with the **Storage Heavy Hitter** dashboard. To understand the IO demand in your environment, use both of them concurrently.

The **Network Top Talkers** dashboard displays sustained hits that last for an hour, as they can cause serious problems in a shared IaaS environment. You can identify the villain VMs and compare their demands with the capabilities of the underlying IaaS.

How to Use the Dashboard

The dashboard shows the current workload. This is the total network load (received and transmitted) from all the vSphere environments monitored by vRealize Operations Manager. The idea is to give you an indicator on how hard the overall load is.

- Select a data center from the data centers list.
 - The columns show the number of clusters, ESXi hosts, and VMs for each data center. The VM count includes the powered off VM. To only see the running VM count, edit the widget.

- If you want to see information from all the data centers, select the vSphere world row.
- Upon selection, the Total Demand Line chart and the Top Talkers tables fill up.
- Total Demand Line Chart
 - The total throughput (received and transmitted) in the selected data center.
 - Displays both, the five minute peak and the hourly average in one line chart. You can click the metric name to hide it.
- Top Talkers Table
 - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important. For example, an ESXi with 2 x10 GB port can theoretically handle 20 GB TX + 20 GB RX as its full duplex.

Points to Note

- Understanding high demand helps you monitor IaaS and plan your capacity. IaaS provides four services, CPU, memory, disk, and network. While CPU, memory, and disk are bound, an active VM can consume all your network bandwidth, packet per second capacity, and the storage IOPS capacity. A VM with 4 vCPU and 16 GB memory cannot consume more than this amount, the same applies to disk space. A VM configured with 100 GB disk space cannot consume more than that.
- Network throughput, disk throughput, and disk IOPS can spike as their physical limits are very high per VM. This means that IaaS has enough capacity for all workloads and performs well until the VMs start consuming abnormally high amounts of network and disk bandwidth.

Storage Heavy Hitters Dashboard

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard. To understand the IO demands in your environment, use both of them together. If you are using ethernet-based storage, storage traffic runs over the same physical network as your ethernet-based network traffic.

Design Considerations

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard, so they share a consideration behind their design. For more information, see [Network Top Talkers Dashboard](#).

How to Use the Dashboard

- See the **Network Top Talkers** dashboard as they have the same design.
 - The main difference between **Storage Heavy Hitters** and **Network Top Talkers** is that the storage IO has two dimensions: IOPS and throughput.
 - Network IO does not have the IOPS dimension as the packet size is identical (1500 bytes being the standard packet, and 9000 bytes being the jumbo frames).

- Storage IOPs and throughput are related, so use both to gain insight, they should display a similar pattern. If not, that indicates varying block sizes. For example, a throughput spike without an accompanying IOPs spike indicates large block sizes.
- Which VMs hit the storage the hardest.
 - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important, because different classes of SSD have different IOPS and throughput capabilities.

After identifying the villain VM, talk to the VM owners if the numbers are excessive during peak hours and identify the reasons behind the excessive usage. You must ensure that they do not create a hot spot. For example, vSAN cluster with > 100 disk can handle numerous IOPS but if the VM objects are only on a few disks, those disks can become a hot spot.

Points to Note

- Interpreting IOPs and throughput metrics depends on your underlying physical storage. For visibility into this hardware layer, add physical storage metrics to the dashboard.

VM Contention Dashboard

The **VM Contention** dashboard is the primary dashboard for VM performance. It is designed for VMware administrators or architects. It can be used for both, monitoring, and troubleshooting. Once you determine that there is a performance issue, use the **VM Utilization** dashboard to see if the contention is caused by high utilization.

Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

For understanding the performance concept of the selected counters and their thresholds, see the [Performance Dashboards](#)

How to Use the Dashboard

- Select a data center from the data center table.
 - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

Note The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- The two bar charts are automatically shown.
 - Use them together to get an insight about your CPU readiness and your Memory contention analysis. Analyze how the cluster serves the VMs. For each VM, it picks the worst metric in the last 24 hours. By default, vRealize Operations Manager collects data every 5 minutes, so this is the highest value among 288 datapoints. Once it has the value from each VM, the bar charts puts each VM in the respective performance buckets. The threshold in the buckets considers best practices, hence they are color coded.
 - For any critical environment, expect that all the VMs are served well by the IaaS. You must see green on both distribution charts. For development purposes, you can tolerate a small amount of contention in both CPU and Memory.
- VM Performance in selected Data Center.
 - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
 - The table is sorted by KPI Breach columns, directing your attention to the VMs that are not served well by the IaaS.
 - The table shows the hostnames known by Windows or Linux. This is the name that the application team or VM owner knows, as they might not be familiar with the VM name.
 - The rest of the columns show performance counters. Because the goal is proactive monitoring, the counters are the worst and not the average, during the monitoring period. Because the operations context here is performance, not capacity, the table considers the last 24 hours only. Daily use is encouraged as any activity older than 24 hours is considered irrelevant from a performance troubleshooting viewpoint.
 - The column KPI Breach counts the number of SLA breaches in any given 5 minutes. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter varies from 0–4, with 0 being the ideal. The value 4 indicates that all 4 IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, not an external SLA. Your internal threshold should be more stringent, so that you have a reaction time.
- Select a VM from the table.
 - All the health charts show the KPI of that VM.
 - The health charts display the last value, lowest value, and the peak value. Expect that the peak is within your threshold.

Points to Note

- This dashboard uses Guest OS counters and VM counters appropriately. The two layers are distinct layers, and they each provide a unique visibility that the other layers might not give.

For example, when the VMkernel de-schedules a VM as it has to process something else (for example, other VM, kernel interrupt). The Guest OS does not know the reason. In fact, it experiences frozen time for that particular vCPU running on the physical core and experiences time jumps when it is scheduled again.

- Guest OS counters logically require VMware Tools.
- The health chart is color coded. Change the settings if it does not suit your environment. If you are unsure of what suitable numbers to set for your environment, profile the metrics. The [Guest OS Performance Profiling Dashboard](#) dashboard provides an example of how to profile metrics.
- For a smaller environment with one or two data centers, change the filter from data center to cluster. Once you are list a cluster, you can then add the cluster performance (%) metric and sort them in an ascending order. This way the cluster that needs immediate attention is on the top.
- If you have a screen real estate, group the VMs by cluster or by ESXi host. This way, you can quickly see if the problem is in a particular cluster or ESXi host.
- Change the default timeline from one week to one day as and when required to suit your operations.
- If you navigate a lot to the **VM Utilization** dashboard from this dashboard, add a connection using the dashboard to dashboard navigation feature. For more details, see [Dashboard Navigation Details](#).

VM Utilization Dashboard

The VMware administrator uses the **VM Utilization** dashboard with the **VM Contention** dashboard for managing performance.

Design Considerations

Use the **VM Utilization** dashboard to identify virtual machines with a high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when a queue develops inside the Windows or Linux operating systems. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- Select a data center from the data center table.
 - For a smaller environment, select vSphere World to see all the VMs from all the data centers.
-
- Note** The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.
-
- VM Peak CPU Usage (%).
 - There is no peak memory use as it is not applicable. Memory is a form of storage, for example consider a hard disk occupied space. A 90% utilization of the total space is not slower than 10%. This means that the issue is related to capacity issue and not performance.
 - The bar chart is color coded using five colors instead of four. The color gray is introduced to convey any wastage. Resources that are hardly utilized do not signify that the performance is at its peak. It can also mean the opposite. For example, if a VM needs 1+ vCPU, configuring it with 2 CPUs results in better performance instead of configuring it with 128 CPUs.
 - VM Peak Utilization.
 - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
 - The table focuses on peak utilization, because the context is performance and not capacity.
 - Select a VM from the table.
 - All the health charts show the KPI of that VM.
 - Complement the free memory with the memory IOPS or the memory throughput metric. The metrics in a gigabyte measure the space, and not the speed. Memory is a form of storage, so what you must measure is the rate, for example, read-write per second.

Points to Note

- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the [VM Contention Dashboard](#).

Troubleshoot an Application Dashboard

The VMware vRealize Application Management Pack provides discovered applications to be managed in vRealize Operations Manager. Using the **Troubleshoot an Application** dashboard, users can see the applications and the relevant metrics and alerts for the selected application. The dashboard also displays its relationship to the infrastructure. In the list of metrics, select a metric to see its trend over time.

Cluster Contention Dashboard

The **Cluster Contention** dashboard is the primary dashboard for vSphere cluster performance. It is designed for VMware administrators or architects. It can be used for both, monitoring and troubleshooting. Once you determine that there is a performance issue, use the **Cluster Utilization** dashboard to see if the contention is caused by high utilization.

Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

Utilization of the cluster is not shown in the **Cluster Contention** dashboard. You must separate the two concepts: utilization and contention. Performance and capacity are different concepts managed by two separate teams. Both CPU and memory are also shown separately. You can have a problem with one, without any issue in the other. CPU is more common as memory tends to have a lower overcommit ratio.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- Average Cluster Performance (%).
 - This is the primary KPI for your entire IaaS. It plots how your IaaS is performing every 5 minutes, giving you the trend view of the overall performance.
 - The metric itself is simply the average of the Cluster KPI / Performance (%) metric. This performance metric in turn averages the VM Performance / Number of KPIs Breached metric from all the running VMs in the cluster. Hence a value of 100% indicates that every running VM in the cluster is served well.
 - As this KPI takes into account every running VM in your environment, the number should be steady. The analogy in real life is the stock market index. While individual stocks can be volatile, overall the index should be relatively steady on a 5 minutes by 5 minutes basis.
 - The relative movement of the metric is as important as the absolute value of the metric. Your absolute number might not be as high you want it to be, but if there are no complaints for a long time, then there is no urgent business justification to improve it.
- Clusters Performance.
 - It lists all the clusters, sorted by the least performing cluster in the last one week. You can change this time period.
 - The worst performance shows the lowest number in the time period. As vRealize Operations Manager collects data every 5 minutes, there are $12 \times 24 \times 7 = 2016$ data points in a week. This column shows the worst point among these 2016 datapoints.

- A single number among 2016 datapoints can be an outlier that needs to be complemented with another number sometimes. A logical choice is the average of these numbers. For the average performance to be low, a lot of criterias have to be low. Waiting for the average causes a delay in your operations, and rise in complaints. For performance monitoring, the 95th percentile is a better summary than the average.
- Your cluster should function at a 100% and perform its fuctions as planned.
- Select a cluster from the table.
 - All the health charts show the KPI of the selected cluster.
 - For performance, it is important to show both the depth and breadth of the performance problems. A problem that impacts one or two VMs requires a different troubleshooting than a problem that impacts all the VMs in the cluster.
 - The depth is shown by reporting the worst among any VM counter. So the highest value of VM CPU Ready, VM Memory contention, and VM Disk Latency among all the running VMs are shown. If the worst number is good, then you do not need to look at the rest of the VMs.
 - A large cluster with thousands of VMs can have a single VM experiencing poor performance while 99.9% of the VM population is fine. The depth counter might not report that most VMs are fine. It only reports the worst. This is where the breadth counters come in.
 - The breadth counters report the percentage of the VM population that is experiencing performance problem. The threshold is set to be stringent, as the goal is to provide early warning and enable proactive operations.

Points to Note

It is possible for VMs in the cluster to suffer from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at individual consumer (VM). The following table shows various possible reasons.

Event	Aware?
Power Management	No
HT	No
Ready	No
Co-Stop	No
System	No
Steal	No
IO Wait	No
Memory Wait	No

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

Certain settings such as DRS automation level and the presence of many resource pools can impact performance. Consider adding a property widget to show the relevant property of a selected cluster, and a relationship widget to show resource pools.

For a large environment with many clusters, add a grouping to make the list more manageable. Group it by class of service, so you can focus more on the critical clusters.

Cluster Utilization Dashboard

The VMware administrator uses the **Cluster Utilization** dashboard with the **Cluster Contention** dashboard for performance management.

Design Considerations

This dashboard supports the **Cluster Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when VMs experience a contention. By default, vRealize Operations Manager has a 5-minutes collection interval. For five minutes, there may be 300 seconds worth of data points. If a spike is experienced for a few seconds, it may not be visible if the remaining of the 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- CPU(%) and Memory (%).
 - Review the CPU and Memory distribution charts for an overview of the CPU and memory utilization of the clusters.
 - The highest metric in the last one week is used. Average or 95th percentile is not used as this is utilization and not contention. High utilization does not mean bad performance.
 - One week is used instead of one day to give you a longer time horizon and covers the weekend. Adjust the timeline as you deem fit for your operations.
 - Expect memory to be higher than CPU, as it is a form of cache. The Memory Consumed counter is used as it is more appropriate than the Memory Active counter.
 - Low utilization can actually indicate bad performance, as not much of real work gets done. The chart uses the dark gray color for low utilization.

- Clusters Utilization.
 - The cluster utilization table lists all the clusters, sorted by the highest utilization in the last one week. If the table displays the green color, then there is no need to analyze further.
 - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a cluster from the table.
 - All the utilization charts show the key utilization metrics of the selected cluster.
 - For memory, the high utilization counters are explicitly shown, Balloon, Compressed, and Swapped. Notice they exist even though utilization is not even at 90%, indicating high pressure in the past. If you look only at utilization, you might think you are safe.
 - The line charts show both average and highest among ESXi hosts in the cluster. The reason is unbalanced and it is not rare. There are many settings that can contribute to it (for example, DRS settings, VM Reservation, VM – Host Affinity, Resource Pool, Stretched Cluster, and Large VMs).
 - The disk IOPS is split into read and write to gain insight into the behavior. Some workload is read oriented, while others are write oriented.
 - The disk throughput is not shown as it sums all the traffic. In reality, each ESXi host has its own limit.
 - The vMotion line chart is added, as a high number of vMotion can indicate that the cluster load is volatile, assuming the DRS Automation level is not set to the most sensitive setting.

Points to Note

- If your operations team have some forms of standardization that utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold.
- Consider adding a third distribution chart. Show the balloon counter in this third chart, as it complements the consumed counter. If there is no ballooning, a high consumed value is in fact better than a lower value.
- The workload metric can exceed a 100% because it is $\text{demand} / \text{usable capacity} * 100$. This can happen if you have four hosts in a cluster with each host running at 100% demand and admission control is set to 50%.
- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the [Cluster Contention Dashboard](#).

VM Rightsizing Dashboard

The **VM Rightsizing** dashboard helps you adjust the VM size for optimal performance and capacity. It covers both undersized and oversized scenarios. This dashboard is designed for the Capacity and the Operations teams, as rightsizing VMs benefit the day-to-day performance.

Design Considerations

The **VM Rightsizing** dashboard helps you visualize information differently by providing choices for customization. It focuses on a summary that helps discussions with senior management. The reclamation size is grouped into buckets so that you can focus on the largest reclamation opportunities first.

How to Use the Dashboard

Select a data center from the **Datacenters** widget.

- The cluster capacity remaining is displayed to give a better context. Focus on reclaiming the cluster that is low on capacity remaining, and on upsizing the cluster with a high capacity remaining.

Once you select a data center from the **Datacenters** widget, all the remaining widgets automatically display the information of the selected data center.

- There are two widgets for upsizing recommendation, one for CPU and one for Memory.
- There are two widgets for downsizing recommendation, one for CPU and one for Memory.
- The business processes for oversized and undersized VMs are different, as one requires the affected VM to be shut down and the Owner to return the resources. To upsize, you must add incrementally. To downsize, you must remove in one change window as the effort to reduce is the same and there will be only one downtime.

Points to Note

- The metrics used are `Summary|Oversized|Virtual CPUs` and `Summary|Undersized|Virtual CPUs`. It stores the capacity engine calculation on the recommended number of vCPUs that must be removed or added.
- When you change the VM configuration, the application setting might have to change. This is especially on applications that manage its memory (for example, database and JVM), and schedules the fixed number of threads.
- Avoid reducing vCPUs from greater than 1 to 1 for Windows. The SMP kernel is activated during the first installation and the performance can degrade on a uni-processor machine.
- You can enable Hot Add on VM, but take note of the impact on NUMA.
- For more details on rightsizing, see [Rightsizing VMs with vRealize Operations](#).

Datastore Performance Dashboard

Use the **Datastore Performance** dashboard to view performance problems related to storage such as high latency, high outstanding IO, and low utilization. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

Design Considerations

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- Select a data center from the data centers table.

- The list of shared datastores in the data center is shown with their KPI.

Note Datastores that are unavailable are not shown.

- Datastore Performance.

- Read and Write latency are shown separately in the datastore performance table for a better insight. The nature of read and write problems might not be the same so it is useful to see the difference.
 - Both the worst (peak) performance and 95th percentile are shown. If the latter is close to the peak and high, then it is a sustained problem. If the latter is low, then it is a short duration.
 - The table is color-coded. If your operations require a different threshold, edit the widget to adjust it accordingly.

- Select a datastore you want to troubleshoot.

- Its read latency, write latency and outstanding IO is automatically shown.

Note The latency is the normalized average of all VMs in the datastore.

- Its IOPS and throughput are also displayed. These line charts are not color-coded as it varies per customer. Edit the widget and add your expected threshold. It makes it easier for the Operations team.
 - The list of VMs is displayed.
- Select a VM you want to troubleshoot.
 - Its read latency and write latency are shown.

Note The number is at the VM level. If you suspect one of the virtual disks has a high latency, use the counter Peak Virtual Disk Read Latency (ms) and Peak Virtual Disk Write Latency (ms).

Points to Note

- The vSphere storage is represented as a datastore. The underlying storage protocol can be files (NFS) or blocks (VMFS). vSAN uses VMFS as its consumption layer as it is unique to vSAN, and has its own monitoring need. Latency can happen when IOPS and throughput are not high. When latency occurs, troubleshooting can take much time.

- You can look at the logs and queue in the various storage stacks (for example, driver) and monitor their performance.
- Datastores that share the same underlying physical array can experience problem at the same time. The underlying array can experience a hot spot on its own, as it is made of independent magnetic disks or SSD.
- The dashboard does not have datastore clusters. If your environment uses it, add a View List to list them, and use this view list to drive the Datastore Performance view list.

ESXi Contention Dashboard

The **ESXi Contention** dashboard is the primary dashboard for managing ESXi host performance. The VMware administrator or architect can use it to monitor and troubleshoot any performance issue. If you determine that there is a performance issue, use the **ESXI Utilization** dashboard to see if the cause for the contention is high utilization.

Design Consideration

The **ESXi Contention** dashboard complements the [Cluster Contention Dashboard](#), and shares the same design consideration.

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed to be used daily, hence the views are set to show data in the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- ESXi CPU Performance and ESXi Memory Performance.
 - Review the two distribution charts for an overview of all the ESXi host's utilization and memory performance.
 - Both charts are using the percentage of VM facing performance counter and not the worst performance among VM counter because you are looking at the ESXi performance and not at the single VM performance. See how it handles all the VMs.
 - The bar chart is color coded. Keep the percentage of the VM population not being served under 10%.
- ESXi Hosts Performance.
 - The ESXi hosts performance table lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table displays the green color, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance greater than 24 hours are likely to be irrelevant.
 - You can change the time period to the period of your interest. The maximum number is reflected accordingly.

- Select an ESXi host from the table.
 - All the health charts show the KPI of the selected cluster.
 - For performance, it is important to show both depth and breadth of a performance problem. A problem that impacts one or two VMs require a different troubleshooting than a problem that impacts all VMs in the cluster.
 - Worst CPU overlap among VMs in the host is included as it indicates a lot of interruptions. A running VM might get interrupted because the VMkernel needs the physical core to run something else. High and frequent numbers of interruptions are not healthy and can impact the VM performance.
 - Expect the network error to be 1% and dropped packet to be 0 most of the times, if not always. If it is not zero, analyze it to see if there are any patterns across all ESXi hosts, and bring it up with your network team.

Points to Note

- Consider adding a third distribution chart and display the CPU co-stop counter in this third chart, as it complements the CPU ready counter. If your environment has relatively slow network and storage IO, you can add IO wait too.
- Unlike the **Cluster Performance** dashboard, there is no average ESXi hosts performance (%) at the vSphere World level. The reason is most ESXi hosts are part of a cluster and monitoring should be done at the cluster level.
- Certain settings such as power management and hyper threading can impact the performance. Consider adding a property widget to show relevant properties of a selected ESXi host.

ESXi Utilization Dashboard

The VMware administrator uses the **ESXi Utilization** dashboard with the **ESXi Contention** dashboard to manage performance.

Design Considerations

The **ESXi Utilization** dashboard supports the **ESXi Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

The dashboard complements the [Cluster Utilization Dashboard](#) dashboard, by providing the extra details. Hence it has a similar layout.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- ESXi Hosts Utilization.
 - It lists all the ESXi hosts, sorted by the highest utilization in the last one week. If the table is all displaying the green color, then there is no need to analyze further.
 - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select an ESXi host from the table.
 - All the utilization charts display the key utilization metrics of the selected cluster.
 - For memory, the high utilization counters are explicitly shown, for example balloon, compressed, or swapped. You might notice they exist even though utilization is not even at 90%, indicating that there was a high pressure in the past. If you look at only utilization, you might think you are safe.
 - The disk IOPS and the disk throughput are split into read and write to gain an insight into the behavior. Some workload is read oriented, while others are write oriented.
 - The network throughput is split into sent (transmit) and received to gain insight into the behavior. The total usage can be misleading because it sums up the send and receive traffic. In reality the network pipe is one for each direction (due to the full duplex nature of Ethernet), and not shared.

Points to Note

If your operations team have some forms of standardization that the utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold. For more information, see the points to note in the [ESXi Contention Dashboard](#).

Network Performance Dashboard

Use the **Network Performance** dashboard to view performance problems related to network such as high latency, frequent retransmit, and many dropped packets. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

Design Considerations

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

The dashboard enables you to drill down from the distributed switch to the ESXi host and port groups in the switch, and then to the VM.

How to Use the Dashboard

- Distributed Switches.
 - The distributed switches table lists all the switches, sorted by the highest packet dropped. The table splits the incoming traffic and the outgoing traffic for better analysis.
 - As the focus is on performance and not capacity, the throughput counters are not shown.
- Select a switch from the distributed switches table.
 - The health chart shows the dropped packet trend over time.
 - It does not narrow down the list of port groups automatically, as the list of port groups are always showing all the port groups in your environment.
 - If necessary, expand the two collapsed widgets. They show the network throughput and broadcast packets. Utilization is also shown so that you can correlate and understand whether the dropped packets are due to higher utilization.
- Port Groups and ESXi Hosts in the selected switch.
 - They get listed when you select a switch from the distributed switches table.
 - Just like the distributed switch, you can also see their relevant counts.
- If your environment has unused network switches, you can filter them out from this list, as this dashboard focuses only on performance.

Points to Note

- vSphere network is by nature distributed. Each ESXi contributes to the physical NIC. This represents the physical capacity. Distributed switch and its port groups span across these independent network cards. This makes it harder to define and measure its performance. An unbalance can happen among ESXi hosts or physical NIC. In a sense, it is like distributed storage (vSAN). Capacity management does not apply to a port group, since its upper limit (also known as the physical capacity) can vary by even a minute.
- Latency within a data center should be below 1 millisecond. Use vRealize Network Insight to study the latency or the retransmitting problems, caused by moving into the lateral traffic.
- Add a physical network using the appropriate management pack.

Most packets are unicast, between a pair of sender and receiver. If your environment has many VMs sending broadcast packets to everyone and multicast packets to many targets, add a Top-N widget to find out which VMs are sending these packets.

vSAN Contention Dashboard

The **vSAN Contention** dashboard is the primary dashboard for managing vSAN performance. The VMware administrator or architect can use it to monitor and troubleshoot the vSAN cluster performance. If you determine that there is a performance issue, use the **vSAN Utilization** dashboard to see if the cause for the contention is high utilization.

Design Considerations

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

The **vSAN Contention** dashboard complements the [ESXi Contention Dashboard](#), and shares the same design consideration. It focuses on the storage and vSAN specific metrics, and does not repeat what is already covered. It does not list any non vSAN cluster.

How to Use the Dashboard

- vSAN Peak VM Latency, vSAN Peak CPU Ready, vSAN Peak Dropped Packet.
 - Review the three distribution charts for an overview of all the vSAN clusters performance.
 - The vSAN peak VM latency chart shows the distribution of disk latency experienced by all the VMs in the cluster. You should expect most of the VMs to experience latency that matches your expectation. For example, in an all flash systems, the VMs should not have >20 ms disk latency. If your vSAN environment is all flash, you must adjust the distribution bucket to a more stringent set.
 - The vSAN peak CPU ready chart shows if any of the vSAN kernel modules has to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not wait for CPU time. vSAN gets higher priority than VM World as it lives in the kernel space.
 - The vSAN peak dropped packet chart shows if any of the vSAN clusters are dropping packet in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.
- vSAN Clusters.
 - It lists all the vSAN clusters, sorted by the least performing.
 - It lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is showing all green, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance issues greater than 24 hours are likely to be irrelevant.
 - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a vSAN cluster from the vSAN clusters table.
 - All the health charts show the KPI of the selected cluster.
 - If you are using SMART, the two heat maps at the bottom of the dashboard provide early warning.

Points to Note

- A large vSAN cluster can have many components. Each of these components can have multiple performance metrics. The total number of KPI can reach hundreds of metrics. For example, take a 10 node cluster. It can have 530 counters to check. vRealize Operations Manager aggregates them by introducing a set of KPIs. This analysis reduces the number to a more manageable number. The following table shows the KPIs and their formula.

Name	What it is
Max Capacity Disk Latency (ms)	Highest latency among all capacity disks take the worst, not average, as the latency in a single capacity disk is already an average of all its VMs. If there are 50 VMs on the disk and 30 are issuing IO on it, then its average is among 30.
Min Disk Group Write Buffer Free (%)	Lowest free capacity among all the disk group write buffers. If this number is low, one of your buffers is not enough. While you want to maximize your cache, a low number is an early warning for capacity management.
Max Disk Group Read Cache/Write Buffer Latency (ms)	Each disk has a Read Cache Read Latency, Read Cache Write Latency (for writing into cache), Write Buffer Write Latency, and Write Buffer Read Latency (for de-staging purpose). This takes the highest among all these four numbers and the highest among all disk groups. It is the max of the max because each of the four datapoints is an average of all the VMs on it.
Sum Disk Group Errors	Sum of the bus reset + sum of commands canceled among all the disk groups. You must use sum and not get the max as each member should return zero.
Count Disk Group Congestion Above 60	The number of disk groups congestion greater than 60. 60 is hardcoded in the vSAN Management Pack as it is a good starting point. As any congestion above 60 serves an early warning, count how many of such occurrences happen.
Max Disk Group Congestion	The highest congestion among all disk groups. A high number indicates that at least one disk group is not performing.
Min Disk Group Capacity Free (%)	The lowest free capacity among all disk groups. A low space triggers rebalance.
Min Disk Group Read Cache Hit Rate (%)	The lowest hit rate among the disk group read cache. Ensure that this number is high as it indicates that the read is served by cache.
Sum vSAN PortGroup Packets Dropped (%)	Sum of all vSAN VMkernel port RX dropped packet + TX dropped packet. You should expect no dropped packet in your vSAN network.

vSAN Utilization Dashboard

The VMware administrator uses the **vSAN Utilization** dashboard with the **vSAN Contention** dashboard to manage performance.

Design Consideration

The **vSAN Utilization** dashboard supports the **vSAN Contention** dashboard. Use it to identify vSAN clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, vRealize Operations Manager has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the [Performance Dashboards](#).

How to Use the Dashboard

- Clusters Utilization.
 - It lists all the vSAN clusters, sorted by the least performing.
- Select a vSAN cluster from the clusters utilization table.
 - All the health charts show the KPI of selected cluster.
- Disk Groups
 - It lists all the vSAN clusters, sorted by the least performing.
- Select a Disk Group from the disk groups table.
 - All the health charts show the KPI of selected cluster.

Points to Note

- The **vSAN Utilization** dashboard complements the **vSAN Contention**. For more information, see the points to note in the [vSAN Contention Dashboard](#).

vSAN File Services

The VMware administrator uses the **vSAN File Services** dashboard to monitor the file services running in their vSAN environment.

Design Considerations

This dashboard is designed to complement the vSAN file services management provided by the vCenter Server. The vCenter Server is more of an administrative tool, while vRealize Operations Manager is more of an operations tool. Each tool performs their specific functions and does not duplicate information.

How to Use the Dashboard

- File Shares by Used Space and Latency.
 - Review the file shares by used space and latency heat map.
 - It shows all the file shares in your environment.
 - The greater the use (consumption), the greater the box, so you can easily see the most consumed ones.
 - The file shares are colored by latency. You must watch out for boxes with red color.
- vSAN Clusters with File Services enabled.
 - It lists all the vSAN clusters with file services enabled, giving a convenient view to see which clusters have these settings turned on.
- Select a vSAN cluster from the vSAN clusters with file services enabled table.
 - The file servers in the selected vSAN cluster are shown. When you select a file server, it filters the file shares list to show the file shares in the selected file server.
 - The file shares in the selected vSAN cluster are shown. Selecting a file share displays all the relevant KPI on the file share.

Points to Note

vSAN File Servers and vSAN File Shares are two new objects in vRealize Operations Management Pack for vSAN.

Dashboard Library

Deprecated Dashboards

Deprecated dashboards are kept intact and are not updated as the changes in the new predefined dashboards are substantial. Deprecated dashboards will be kept for at least one release. See the Release Notes for information about why the dashboards are deprecated.

Capacity Allocation Overview Dashboard

This dashboard provides an overview of allocation ratios for virtual machines, vCPUs, and memory for a specific data center or cluster.

Cluster Configuration Dashboard

The Cluster Configuration dashboard provides a quick overview of your vSphere cluster configurations. The dashboard highlights the areas that are important in delivering performance and availability to your virtual machines. The dashboard also highlights if there are clusters which are not configured for DRS, High Availability (HA), or admission control to avoid any resource bottlenecks or availability issues when a host fails.

The heat map in this dashboard helps you to identify if you have hosts where vMotion was not enabled as this may not allow the VMs to move from or to that host. This may cause potential performance issues for the VMs on that host if the host gets too busy. You can also view how consistently your clusters are sized and whether the hosts on each of those clusters are consistently configured.

The Cluster Properties widget in this dashboard allows you to report on all these parameters by exporting the data. You can share the data with the relevant stakeholders within your organization.

You can use the dashboard widgets in several ways.

- **vSphere DRS Status, vSphere HA Status, and HA Admission Control Status:** Use these widgets to view if there are clusters that are not configured for DRS, HA, or admission control. With the information, you can avoid resource bottlenecks or availability issues when a host fails.
- **Is vMotion enabled on hosts in a cluster:** Use this widget to identify if you have hosts where vMotion was not enabled. If vMotion is not enabled, the VMs do not move from or to the host and causes potential performance issues in the VMs on that host if the host gets too busy.
- **Host Count across Clusters:** Use this widget to view all the clusters in your environment. If the clusters have a consistent number of hosts, the boxes displayed are of equal size. This representation helps you determine whether there is a large deviation among cluster sizes, whether there is a small cluster with fewer than four hosts, or whether there is a large cluster. Operationally, keep your clusters consistent and of moderate size.
- **Attributes of ESXi Hosts in the Selected Cluster:** Use this widget to view the configuration details for the hosts within a cluster.
- **All Clusters Properties:** Use this widget to view the properties for all the clusters in the widget.

Cluster Utilization Dashboard

The Cluster Utilization dashboard helps you identify vSphere clusters that are extensively consumed from a CPU, memory, disk, and network perspective.

You can use this dashboard to identify the clusters that cannot serve the virtual machine demand.

You can select a cluster with high CPU, memory, disk, or network demand. The dashboard lists the ESXi hosts that are a part of the given cluster. If there is an imbalance in the use of hosts within the selected clusters, you can balance the hosts by moving the VMs within the cluster.

You can use this dashboard to view the historical cluster demand. If the situation is critical, use Workload Balance and move the VMs out of the clusters to avoid potential performance issues. For more information, see [Configuring and Using Workload Optimization](#). If all the clusters in a given environment display the same pattern, you might have to add new capacity to cater to the increase in demand.

Datastore Usage Overview Dashboard

The Datastore Usage Overview dashboard provides a view of all the virtual machines in your environment in a heat map. The dashboard is suitable for an NOC environment.

The heat map contains a box for each virtual machine in your environment. You can identify the virtual machines that are generating excessive IOPS because the boxes are sized by the number of IOPS they generate.

The colors of the boxes represent the latency experienced by the virtual machines from the underlying storage. An NOC administrator can investigate the cause of this latency and resolve it to avoid potential performance problems.

Datastore Utilization Dashboard

The Datastore Utilization dashboard helps you identify storage provisioning and utilization patterns in a virtual infrastructure.

As a best practice, ensure that the datastores are of standard size, to manage storage in your virtual environments. The heat map on this dashboard displays all the datastores monitored by vRealize Operations Manager and groups them by clusters.

The dashboard uses colors to depict the utilization pattern of the datastores. Grey represents an underutilized datastore, red represents a datastore that has run out of disk space, and green represents an optimally used datastore. You can select a datastore from the dashboard to see the past utilization trends and forecasted usage. The dashboard lists all the VMs that run on the selected datastore. You can reclaim storage used by large VM snapshots or powered off VMs.

You can use the vRealize Operations Manager action framework to reclaim resources by deleting the snapshots or unwanted powered off VMs.

- **Datastore Capacity and Utilization:** Use this widget to find out which datastores are overused and which ones are underused. You can also find out whether the datastores are of equal size. When you select a datastore from this widget, the dashboard is automatically populated with the relevant data.
- **VMs in the Selected Datastore:** Use this widget to view a list of VMs based on the datastore you select. You can also view relevant details such as whether the VMs are powered on and the size of the snapshot if any.
- **Usage Trend of Selected Datastore:** Use this widget to find out the trends in capacity used by a selected datastore as against the total capacity available.
- **All Shared Datastores in the Environment:** Use this widget to view a list of datastores that are shared in your environment. The information displayed in this widget helps you make an informed decision about whether you have to rebalance the capacity of the datastores based on usage.

Distributed Switch Configuration Dashboard

The Distributed Switch Configuration dashboard allows you to view details of virtual switch configuration and utilization. When you select a virtual switch, you can see the list of ESXi hosts,

distributed port groups, and virtual machines that use or are on the selected switch. You can also find out which ESXi hosts and VMs use a specific switch.

You can identify misconfigurations within various network components by reviewing the properties listed in the views within the dashboard. You can track important information such as the IP address and the MAC address assigned to the virtual machines.

As a network administrator, you can use this dashboard to get visibility into the virtual infrastructure network configuration.

You can use the dashboard widgets in several ways.

- **Select a Distributed Switch:** Use this widget to select the switch for which you want to view details. You can use the filter to narrow your list based on several parameters. After you identify the switch that you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Distributed Port Groups on the Switch:** Use this widget to view the port groups on the switch, how many ports each switch has, and the usage details.
- **ESXi Hosts/VMs Using the Selected Switch:** Use these widgets to find out which ESXi hosts and VMs use the selected switch. You can also view configuration details about the ESXi hosts and VMs that use the selected switch.

Heavy Hitter VMs

The Heavy Hitter VMs dashboard helps you identify virtual machines which are consistently consuming a large amount of resources from your virtual infrastructure. In heavily over-provisioned environments, this might create resource bottlenecks resulting in potential performance issues.

You can use this dashboard to identify the resource utilization trends of each of your vSphere clusters. With the utilization trends, you can also view a list of VMs within those clusters based on their resource demands from the CPU, memory, disk, and network within your environment. You can also analyze the workload pattern of these VMs over the past week to identify heavy hitter VMs which might be running a sustained, heavy workload that is measured over a day, or bursty workloads that is measured using peak demand.

You can export a list of offenders and take appropriate action to distribute this demand and reduce potential bottlenecks.

You can use the dashboard widgets in several ways.

- **Select a Cluster:** Use this widget to select a cluster. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Cluster CPU and Cluster Memory:** Use these widgets to view the CPU and memory for the cluster.
- **Cluster IOPS and Cluster Network Throughput:** Use these widgets to view the IOPS and network throughput for the cluster.

- Use the other widgets in the dashboard to view which VMs in the cluster generated the highest network throughput and IOPS. You can also view which VMs in the cluster generated the highest CPU demand and the highest memory demand. You can compare the information for the VM with the results for the cluster and correlate the trends. You can manually set the time to the time period for which you want to view data.

Host Configuration Dashboard

The Host Configuration dashboard provides an overview of your ESXi host configurations, and displays inconsistencies so that you can take corrective action.

The dashboard also measures the ESXi hosts against the vSphere best practices and indicates deviations that can impact the performance or availability of your virtual infrastructure. Although you can view this type of data in other dashboards, in this dashboard you can export the ESXi configuration view and share it with other administrators.

Host Usage Overview Dashboard

The Host Usage Overview dashboard provides a view of all the ESXi hosts in your environment in a heat map. The dashboard is suitable for an NOC environment.

Using this dashboard an NOC administrator can easily find resource bottlenecks created due to excessive Memory Demand, Memory Consumption or CPU Demand.

The heat map displays hosts grouped by clusters to help you locate clusters that are using excessive CPU or memory. You can also identify if you have ESXi hosts within the clusters that are not evenly utilized. An administrator can then trigger activities such as workload balance or set DRS to ensure that hot spots are eliminated.

Host Utilization Dashboard

The Host Utilization dashboard helps you identify hosts that are extensively consumed from a CPU, memory, disk, and network perspective.

You can use this dashboard to identify hosts that cannot serve the virtual machine demand. The dashboard provides a list of the top 10 virtual machines. You can identify the source of this unexpected demand and take appropriate actions.

You can use the dashboard to view demand patterns over the last 24 hours and identify hosts that have a history of high demand. You must move the virtual machines out of these hosts to avoid potential performance issues. If all the hosts of a given cluster display the same pattern, you might have to add new capacity to cater to the increase in demand.

Migrate to vSAN

The Migrate to vSAN dashboard provides you with an easy way to move virtual machines from existing storage to newly deployed vSAN storage.

You can use this dashboard to select non-vSAN datastores that might not serve the virtual machine IO demand. By selecting the virtual machines on a given datastore, you can identify the historical IO demand and the latency trends of a given virtual machine. You can then find a suitable vSAN datastore which has the space and the performance characteristics to serve the demand of this VM. You can move the virtual machine from the existing non-vSAN datastore to the vSAN datastore. You can continue to watch the use patterns to see how the VM is served by vSAN after you move the VM.

Operations Overview Dashboard

The Operations Overview dashboard provides you with a high-level view of objects which make up your virtual environment. You can view an aggregate of the virtual machine growth trends across the different data centers that vRealize Operations Manager monitors.

You can also view a list of all your data centers with inventory information about how many clusters, hosts, and virtual machines you are running in each of your data centers. By selecting a particular data center, you can narrow down on the areas of availability and performance. The dashboard provides a trend of known issues in each of your data centers based on the alerts which have triggered in the past.

You can also view a list of the top 15 virtual machines in the selected data center which might be contending for resources.

You can use the dashboard widgets in several ways.

- **Environment Summary:** Use this widget to view a summary of the overall inventory of your environment.
- **Select a Datacenter:** Use this widget to select the data center for which you want to view operational information. You can use the filter to narrow your list based on several parameters. After you identify the data center you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Cumulative Up-time of all Clusters:** Use this widget to view the overall health of the clusters in the data center you selected. The metric value is calculated based on the uptime of each ESXi host, when you take into account one host as the HA host. If the number displayed is less than 100%, it means that at least two hosts within the cluster were not operational for that period.
- **Alert Volume (in selected DC):** Use this widget to view the breakdown of alert trends based on their criticality.
- **Top-N:** You can also view a list of 15 VMs that had the highest average CPU contention, the highest use of memory, and the highest disk latency for the last 24 hours. To obtain specific data, you can manually set the time to the time of the problem. To set the time, click the **Edit Widget** icon from the title bar of the widget and edit the **Period Length** drop-down menu.

Optimization History Dashboard

The Optimization History dashboard displays the results of optimization activity.

The Optimization History dashboard belongs to the Optimize group of dashboards. The dashboard covers three optimization benefits; optimize performance, optimize capacity, and optimize virtual machine placement.

Optimizing performance can be performed in vRealize Operations Manager using Workload Optimization, or started on demand. The charts on this row show a box for each data center or custom data center and the optimization recommendation. Green indicates an optimized data center or custom data center. A red box means that optimization might be required, and a white box means that optimization is not configured for that object.

For capacity optimization, this row provides a summary of the average VM cost per month, the savings that can be achieved through reclaiming idle or powered off virtual machines, or deleting old snapshots.

Virtual Machine Happiness is a term used to describe VMs that are getting the resources they need, when they need them. You can also see recent vMotion activity related to vSphere's Distributed Resource Scheduler, which together with vRealize Operations predictive DRS feature makes sure your VMs are getting the resources they need. Workload placement vMotions are also shown as Non-DRS Moves in the graph.

Optimize Performance Dashboard

The Optimize Performance dashboard helps you identify virtual machines that can be configured to improve overall performance.

The capacity analytics engine intelligently calculates the settings for CPU and memory for virtual machines to give you the best performance and accurate resource allocation for all workloads.

The dashboard organizes virtual machines by undersized - or virtual machines that are not being served well - and oversized - which are virtual machines that are not using all allocated resources. Both categories consider CPU and memory usage and provide recommendations for optimal sizing.

Troubleshoot a Cluster

The Troubleshoot a Cluster dashboard allows you to identify clusters that have issues and isolate them easily.

You can use the search option to identify a cluster that has an issue. You can also sort the clusters based on the number of active alerts.

After you select the cluster you want to work with, you can view a quick summary of the number of hosts in that cluster and the VMs served by the cluster. The dashboard provides you with current and past utilization trends and also known issues in the cluster in the form of alerts.

You can view the hierarchy of objects related to the cluster and review the status to identify if the objects are impacted because of the current health of the cluster. You can quickly identify any contention issues by looking at the maximum and average contention faced by the VMs on the selected cluster. You can narrow down and view those VMs that have resource contention and take specific steps to troubleshoot and resolve issues.

You can use the dashboard widgets in several ways.

- **Search for a cluster:** Use this widget to select the cluster for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the cluster you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Is your cluster busy?:** Use this widget to view the CPU and memory demand.
- **Are there active alerts on your cluster:** Use this widget to view only the critical alerts.
- **Are the relatives healthy?:** Use this widget to view the hierarchy of the objects related to the cluster and if any of the objects are impacted.
- View the maximum and average CPU, memory, and disk latency for the VMs. If the VM faces contention, it might mean that the underlying infrastructure does not have enough resources to meet the needs of the VMs.
- View a list of VMs that face CPU, memory, and disk latency contention. You can then troubleshoot and take steps to resolve the problem.

Troubleshoot a Datastore

The Troubleshoot a Datastore dashboard allows you to identify storage issues and act on them.

You can use the search option to identify a datastore that has an issue or you can identify a datastore that has high latency as seen in red on the heat map. You can also sort all the datastores with active alerts and troubleshoot the datastore with known issues.

You can select a datastore to see its current capacity and utilization with the number of VMs served by that datastore. The metric charts help you view historical trends of key storage metrics such as latency, outstanding IOs, and throughput.

The dashboard also lists the VMs served by the selected datastore and helps you analyze the utilization and performance trends of those VMs. You can migrate the VMs to other datastores to even out the IO load.

You can use the dashboard widgets in several ways.

- **Search for a datastore:** Use this widget to select the datastore for which you want to view performance details. You can use the filter to narrow your list based on several parameters. After you identify the datastore you want to view, select it. The dashboard is automatically populated with the relevant data.
- **Are there active alerts on your datastore:** Use this widget to view only the critical alerts.
- **Are the relatives healthy?:** Use this widget to view the hierarchy of the objects related to the datastore and if any of the objects are impacted.
- **Is your datastore experiencing high latency? and Any outstanding disk I/Os?:** Use these widgets to view those datastores with high latency and outstanding disk I/O trends. Ideally, your datastores must not have outstanding disk I/O.

- **How many IOPS is your datastore serving and Latency trend for the I/Os done by the VM:** Use these widgets to view the current IOPS and latency of the VMs in the selected datastore.
- Use the other widgets in the dashboard to view trends for the selected datastore regarding disk latency, IOPS, and throughput, VMs served by the datastore and I/O pattern of the selected VM.

Troubleshoot a Host

The Troubleshoot a Host dashboard allows you to search for specific hosts or sort hosts with active alerts. ESXi hosts are the main source of providing resources to a VM and are critical for performance and availability.

To view the key properties of each host, select a host from the dashboard. You can ensure that the host is configured according to the virtual infrastructure design. Any deviation from standards might cause potential issues. You can use the dashboard to answer key questions about current and past utilization and workload trends over the last week. You can also view if the VMs served by the host are healthy.

Since the dashboard lists all the critical events that might affect the availability of the hosts, you can view hardware faults associated with the host. You can view a list of the top 10 VMs that demand CPU and memory resources from the identified host.

Troubleshoot a VM Dashboard

The Troubleshoot a VM dashboard helps an administrator to troubleshoot everyday issues in a virtual infrastructure. While most of the IT issues in an organization are reported at the application layer, you can use the guided workflow in this dashboard to help investigate an ongoing or a suspected issue with the VMs supporting the impacted applications.

You can search for a VM by its name or you can sort the list of VMs with active alerts on them to start your troubleshooting process. When you select a VM, you can view its key properties to ensure that the VM is configured as per your virtual infrastructure design. Any deviation from standards may cause potential issues. You can view known alerts and the workload trend of the VM over the past week. You can also view if any of the resources serving the virtual machine have an ongoing issue.

The next step in the troubleshooting process allows you to eliminate the major symptoms which might impact the performance or availability of a VM. You can use key metrics to find out if the utilization patterns of the VMs are abnormal or if the VM is contending for basic resources such as CPU, memory, or disk.

You can use the dashboard widgets in several ways.

- **Search for a VM:** Use this widget to view all the VMs in the environment. You can select the VM you want to troubleshoot. You can use the filter to narrow your list based on several parameters, such as name, folder name, associated tag, host, or vCenter Server. After you identify the VM you want to troubleshoot, select it. The dashboard is automatically populated with the relevant data.

- **About the VM:** Use this widget to understand the context of the VM. This widget also lends insights to analyze the root cause of the problem or potential mitigations.
- **Are there active alerts on the VM?:** Use this widget to view active alerts. To see noncritical alerts, click the VM object.
- **Is the VM working hard over the last week?:** Use this widget to view the workload trend of the VM for the last week.
- **Are the relatives healthy?:** Use this widget to view the ESXi host where the VM is now running. This host might not be the ESXi host where the VM was running in the past. You can view the remaining related objects and see whether they might contribute to the problem.
- **Is the VMs demand spiking or abnormal?:** Use this widget to identify spikes in the VM demand for any of the resources such as CPU, memory, and network. Spikes in the demand might indicate an abnormal behavior of the VM or that the VM is undersized. The memory utilization is based on the Guest OS metric. It requires VMware Tools 10.0.0 or later and vSphere 6 Update 1 or later. If you do not have these products, the metric remains blank.
- **Is the VM facing contention?:** Use this widget to identify whether the VM is facing contention. If the VM is facing contention, the underlying infrastructure might not have enough resources to meet the needs of the VM.
- **Does the cluster serving the VM have contention?:** Use this widget to view the trend for the maximum CPU contention for a VM within the cluster. The trend might indicate a constant contention within the cluster. If there is contention, you must troubleshoot the cluster as the problem is no longer with the VM.
- **Does the datastore serving the VM have latency?:** Use this widget to help you correlate the latency at the datastore level with the total latency of the VM. If the VM has latency spikes, but the datastore does not have such spikes, it might indicate a problem with the VM. If the datastore faces latency as well, you can troubleshoot to find out why the datastore has these spikes.
- **Parent Host and Parent Cluster:** Use these widgets to view the host and the cluster on which the VM resides.

Troubleshoot vSAN Dashboard

The Troubleshoot vSAN dashboard helps you view the properties of your vSAN cluster and the active alerts on the cluster components. The cluster components include hosts, disk groups, or the vSAN datastores.

You can select a cluster from the dashboard and then list all the known problems with the objects associated with the cluster. The objects include clusters, datastores, disk groups, physical disks, and VMs served by the selected vSAN cluster.

You can view the key use and performance metrics from the dashboard. You can also view the usage and performance trend of the cluster for the last 24 hours. You can also view historical issues and analyze the host, disk group, or physical disk.

You can use the heat maps within the dashboard to answer questions about write buffer usage, cache hit ratio, and host configurations. You can also use the heat maps to answer questions about physical issues with capacity and cache disks, such as drive wear out, drive temperature, and read-write errors.

You can use the dashboard widgets in several ways.

- **Search for a vSAN cluster:** Use this widget to search vSAN clusters. You can view the details of each vSAN cluster including the number of hosts, VMs, cache disks, capacity disks, and cluster type are provided. You can also view if the vSAN cluster is dedupe and compression enabled, and stretched.
- **Any alerts on the cluster, hosts, VMs or disks?:** Use this widget to view alerts on the cluster, VMs, or disks in your environment.
- **Are the relatives healthy?:** Use this widget to view the health, risk, and efficiency of the relatives. This widget also allows you to view the health of the datastore in a host and disks in each disk group.
- **Are outstanding I/Os high?:** Use this widget to view the key performance metrics. The widget indicates outstanding I/Os within 24 hours time period.
- **Are VMs facing read latency?:** Use this widget to view the read latency of VMs.
- **Are VMs facing write latency?:** Use this widget to view the write latency of VMs.
- **Is the write buffer low?:** Use this widget to view the usage of the write buffer on diskgroups in a cluster.
- **Are the hosts consistently configured?:** Use this widget to view the participating hosts in the selected cluster and to determine if the hosts are consistently configured.
- **Cache Disks: Any hardware issues?:** Use this widget to view the individual cache disks measured against various metrics.
- **Capacity Disks: Any hardware issues?:** Use this widget to view the individual capacity disks measured against various metrics.

Troubleshoot with Logs Dashboard

When vRealize Operations Manager is integrated with vRealize Log Insight, you can access the custom dashboards and content pack dashboards from the Troubleshoot with Logs dashboard. You can view graphs of log events in your environment, or create custom sets of widgets to access the information that matters most to you.

You can investigate an ongoing issue within your virtual infrastructure using the logs. You can view predefined views created within vRealize Log Insight to answer questions from predefined queries within vRealize Log Insight.

You can correlate metrics and queries within vRealize Operations Manager to troubleshoot issues across applications and infrastructure.

For more information about the Troubleshoot with Logs dashboard, see the [vRealize Log Insight documentation](#).

To access the Troubleshoot with Logs dashboard from vRealize Operations Manager, you must either:

- Configure the vRealize Log Insight adapter from the vRealize Operations Manager interface, or
- Configure vRealize Operations Manager in vRealize Log Insight.

For more information on configuring, see [Configuring vRealize Log Insight with vRealize Operations Manager](#).

Utilization Overview Dashboard

The Utilization Overview dashboard helps you view the available capacity in the virtual infrastructure.

The Utilization Overview dashboard allows you to assess the utilization at each resource group level such as vCenter, data center, custom data center, or vSphere cluster. You can quickly select an object and view the total capacity, used capacity, and usable capacity of the object to understand the current capacity situation.

You can use the dashboard widgets in several ways.

- **Total Environment Summary:** Use this widget to view the total available capacity in the environment including information about the number of hosts and datastores. You can also view storage, memory, and CPU capacity, and the number of physical CPUs.
- **Select an Environment:** Use this widget to select a data center, a cluster compute resource, or a vCenter Server. You can use the filter to narrow your list based on several parameters. After you identify the data center you want to view, select it. The dashboard is populated with the relevant data.
- **Inventory:** Use this widget to view the number of running VMs and hosts. You can also view the number of datastores and the consolidation ratio in the environment.
- **Usable Capacity (Exclude HA Buffers):** Use this widget to view the capacity that is available in the virtual infrastructure.
- **Used Capacity:** Used this widget to view how the capacity is used in various data centers and clusters.
- **Capacity Remaining:** Use this widget to view the capacity remaining in terms of memory, storage, and CPU capacity remaining.
- **Predicted Time Remaining:** Use this widget to view the predicted time remaining based on the use patterns in the environment.
- **Cluster Capacity Details:** Use this widget to view detailed capacity information for each cluster.

VM Configuration Dashboard

The VM dashboard focuses on highlighting the key configurations of the virtual machines in your environment. You can use this dashboard to find inconsistencies in configuration within your virtual machines and take quick remedial measures. You can safeguard the applications which are hosted on these virtual machines by avoiding potential issues due to misconfigurations.

Some of the basic problems the dashboard focuses on includes identifying VMs running on older VMware tools versions, VMware tools not running, or virtual machines running on large disk snapshots. VMs with such symptoms can lead to potential performance issues and hence it is important that you ensure that they do not deviate from the defined standards. This dashboard includes a predefined Virtual Machine Inventory Summary report which you can use to report the configurations highlighted in this dashboard for quick remediation.

You can use the dashboard widgets in several ways.

- Use the Large VMs widgets to view graphical representations of VMs that have a large CPU, RAM, and disk space.
- **Guest OS Distribution:** Use this widget to view a break up of the different flavors of operating systems you are running.
- **Guest Tools Version** and **Guest Tools Status:** Use these widgets to identify if you have inconsistent or older version of VMware tools which might lead to performance issues.
- View the VMs with limits, large snapshots, orphaned VMs, VMs with more than one NIC, and VMs with a nonstandard operating system. These VMs have a performance impact on the rest of the VMs in your environment even though they do not fully use their allocated resources.

You can customize the views in the widgets.

- 1 Click the **Edit Widget** icon from title bar of the widget. The **Edit** widget dialog box is displayed.
- 2 From the **Views** section, click the **Edit View** icon. The **Edit View** dialog box is displayed.
- 3 Click the **Presentation** option in the left pane and make the required modifications.

VM Utilization Dashboard

The VM Utilization dashboard helps you as an administrator to capture the utilization trends of any VM in your environment. You can list the key properties of a VM and the resource utilization trends for a specific time period. You can share the details with the VM or application owners.

The dashboard displays resource utilization trends so that the VM or application owners can view these trends when they expect a high load on applications. For example, activities like batch jobs, backup schedules, and load testing. Application owners must ensure that the VMs do not consume 100% of the provisioned resources during these periods. Excessive consumption of the provisioned resources can lead to resource contention within the applications and can cause performance issues.

- **Search for a VM to Report its Usage:** Use this widget to select the VM you want to troubleshoot. You can use the filter to narrow your list based on several parameters. After you identify the VM that you want to view, select it. The dashboard is automatically populated with the relevant data.
- **About the VM:** Use this widget to view the VM you selected and its details. You select the VM in the Search for a VM to Report its Usage widget.
- **VM Utilization Trend: CPU, Memory, IOPS, Network:** Use this widget to view information about the utilization and allocation trends for CPU demand, memory workload, disk commands per second, and the network usage rate.

vSAN Capacity Overview

The vSAN Capacity Overview dashboard provides an overview of vSAN storage capacity and savings achieved by enabling deduplication and compression across all vSAN clusters.

You can view current and historical use trends, and future procurement requirements from the dashboard. You can view details such as capacity remaining, time remaining, and storage reclamation opportunities to make effective capacity management decisions.

You can view the distribution of use among vSAN disks from the dashboard. You can view these details either as an aggregate or at an individual cluster level.

vSAN Operations Overview

The vSAN Operations Overview dashboard provides an aggregated view of the health and performance of your vSAN clusters.

You can use this dashboard to get a complete view of your vSAN environment and what components make up the environment. You can also view the growth trend of virtual machines served by vSAN.

You can use the dashboard to understand the utilization and performance patterns for each of your vSAN clusters by selecting one from the list that is provided. You can use this dashboard to track vSAN properties such as hybrid or all flash, deduplication and compression, or a stretched vSAN cluster.

You can view the historic performance, utilization, growth trends, and events related to vSAN, with the current state.

You can identify the vSAN encryption status at cluster levels.

vSphere Security Compliance Dashboard

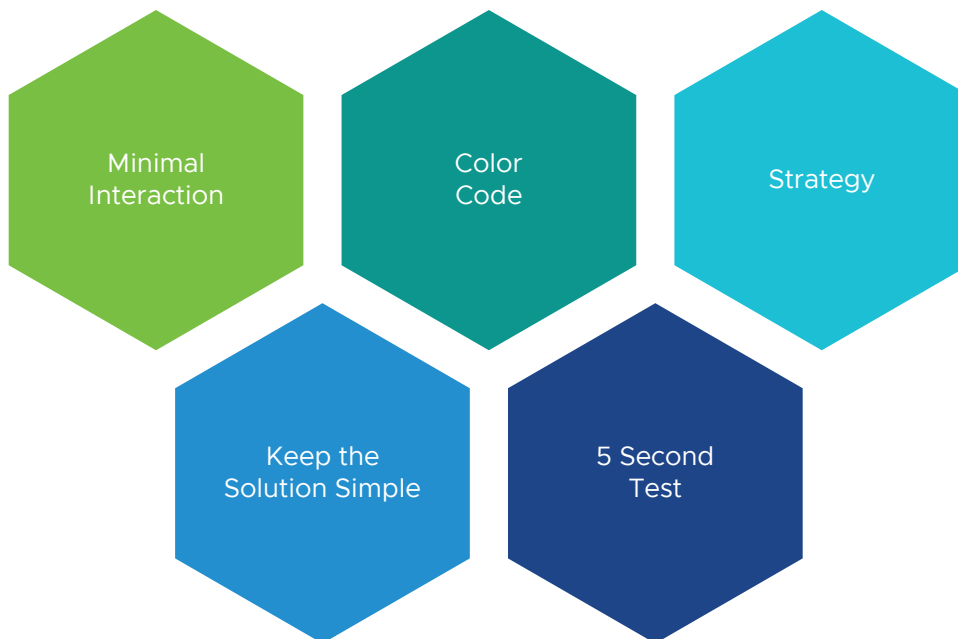
The vSphere Security Compliance dashboard measures your environment against the *vSphere Hardening Guide* and lists any objects which are non-compliant.

This dashboard displays the trend of high risk, medium risk, and low risk violations and shows the overall compliance score of your virtual infrastructure. Using heat maps, you can investigate various components to check the compliance for your ESXi hosts, clusters, port groups, and virtual machines. Each non-compliant object is listed in the dashboard with recommendations on the remediation required to secure your environment.

Executive Summary Dashboards

The requirements of the CIO, Head of Global Infrastructure, and IT Senior Management vary from the requirements of the technical teams. The **Executive Summary** dashboards provide an overall information on capacity and inventory in business terms.

These dashboards allow you to display problems related to budget and resource, and provide visibility to the senior management into the live environment. By doing this, you can prove the need for additional hardware. If there is wastage that has to be reclaimed, you can display where and how large the wastage is using these dashboards. vRealize Operations Manager provides two example dashboards to get you started. As each executive might have a unique requirement or preference, the dashboards can be customized accordingly. The five principles displayed in the following figure are used to design the **Executive Summary**



dashboards.

- Keeping the interaction, such as clicking, zooming, and sorting to a minimal.
- Use of color codes to have a user interface that is easy to understand.
- Each dashboard answers a specific question and the information is presented in business terms.

- Keep the solution simple and have a portal that is easy to access.
- Ensure that the dashboards are understood within five seconds.

Capacity Summary Dashboard

The **Capacity Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Inventory Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

Design Considerations

See [Executive Summary Dashboards](#) for common design considerations among all the dashboards for the IT senior management.

How to Use the Dashboard

The **Capacity Summary** dashboard has two sections:

- The top section of the dashboard provides a summary at the vSphere World level.
 - The **VM Growth** widget displays the weekly average of the VM growth and provides holistic visibility of overall growth across all data centers for both running and powered off workloads. If an increase in the VM count is not accompanied by a corresponding increase in utilization, these newly provisioned VMs are likely not yet used.
 - The **Overcommit Ratio** widget highlights the efficiency gained by vSphere virtualization running multiple workloads on a shared infrastructure. Overcommitment has to be further reviewed along with elevated resource contention to understand the impact of performance on VMs competing for resources. In general, Overcommit is required to be financially more economical than the public cloud. As a reference, AWS typically overcommits CPU 2:1 by counting the hyper-threading and does not overcommit memory.

Note vRealize Operations Manager uses physical CPU Cores not Logical Cores (Hyper-threading) for all CPU-based capacity calculations.

- The bottom section of the dashboard enables drill down into individual compute or storage capacity.
 - Capacity is split into Compute (vSphere Clusters) and Storage (Datastores) views. The heat map displays capacity by size and color by time remaining. By selecting either Clusters or Datastores, you can further drill down understand the remaining capacity and time (in days).

Points to Note

- Capacity remaining is not displayed at the vSphere World level as it can be misleading, especially in a global, or large infrastructure. Clusters also tend to serve a different purpose and they are not interchangeable.

- If you are using both on-prem and external cloud, for example, VMware on AWS, consider splitting the dashboard into two columns.

Inventory Summary Dashboard

The **Inventory Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Capacity Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

Design Considerations

See [Executive Summary Dashboards](#) for common design considerations among all the dashboards for the IT senior management.

How to Use the Dashboard

- The **Summary** widget provides a quick view of the key inventory number.
 - The scoreboard is interactive. This widget drives the eight pie charts that are placed at the bottom of the dashboard. Since all the information is at the vSphere World level, clicking any of them will display details of the total inventory.
- Select any data center from the **Datacenters** widget.
 - This widget drives Clusters and Datastores so that you can quickly view what you have in a given data center and related capacity.
 - For a small environment, the vSphere World is displayed so you can view all the VMs in the environment.
 - To sort by any of the columns in the table, click the column title.
- The eight charts in the dashboard provide details of the inventory. They are driven by **Datacenters**, **Compute**, **Storage**, and **Summary** widgets.

Points to Note

- Understand the relationship hierarchy in vSphere. For example, Compute (cluster) is not a parent of Storage (datastore), so logically it is not possible to display datastores in a cluster. Data center consists of compute (cluster), network (distributed switch), and storage (datastore).
- Datastores do not drive the pie chart. This is a known limitation in the View widget.
- If your senior management wants to view the largest VM in a given environment, add a Top-N widget to list the top 10 largest consumers so that CPU, memory, and the disk details are highlighted.

Network Operation Center

A dashboard projected on the large screen serves a different business purpose than a dashboard on your laptop or desktop. It is placed strategically because it displays a time sensitive

information. Dashboards complement alerts and cannot replace it. The five principles displayed in the following figure are used to design the predefined **Network Operation Center** dashboards.



- Keeping the interaction, such as clicking, zooming, and sorting to a minimal. Avoid having buttons, use of mouse or keyboard to view data.
- Use of color codes to have a user interface that is easy to understand.
- Displaying content that drives action. Display of live information as the focus is on immediate remediation. Problems that need immediate actions are displayed, for example, stop provisioning of new VM or take action on VMs that abuse the shared infrastructure.
- Display of problems that do not require immediate attention are avoided, for example, increase supply of infrastructure, such as adding hardware.
- Keep the display simple and have a portal that is easy to access.
- Dashboards are designed to display minimal and critical information only.
- Displays of numbers in percentage, with 0% being poor and 100% being perfect. To display utilisation, you can use the following markers:
 - 50% indicates good and balanced utilization. However, the ideal value is 75%
 - 0% indicates wastage

- 100% indicates high utilization
- Ensure that the dashboards are understood within five seconds.

Live! Cluster Performance Dashboard

The **Live! Cluster Performance** dashboard provides live information on whether the requests of the VMs are met by their underlying compute clusters. This dashboard focuses on CPU, Memory, and the performance of the clusters. Use this dashboard to view if there is any problem in meeting the demands of the VMs and if there is any unbalance within a cluster. The **Live! Cluster Performance** dashboard is the primary dashboard and it complements the **Live! Cluster Performance** dashboard which is the secondary dashboard. This secondary dashboard displays if the performance problem is caused by high utilization. The primary dashboard answers the question 'Is our IaaS performing?', while the secondary dashboard answers the question 'Is our IaaS working hard?'.

Design Consideration

The **Live! Cluster Performance** dashboard displays three heat maps. The heat maps complement each other and must be used together. The location of each cluster and ESXi hosts within those clusters is identical in all heat maps. The fixed positioning allows you to compare if the problem is caused by memory contention, CPU ready, or CPU co-stop.

The sizes of each cluster and ESXi hosts are constant. Variable sizing creates a distraction and can result in small boxes, making it difficult to read.

The focus of the performance is on the population and not on a single VM. This is not a single VM troubleshooting dashboard but a dashboard focusing on infra problem. As the infra counter is mathematically an aggregation of VM counters, you must have a right roll-up strategy. As the goal is to provide an early warning, do not use the average as a roll-up technique. Use the percentage of the population exceeding a threshold. The threshold is set to be stringent to receive an early warning.

How to Use the Dashboard

Review the heat maps, **Memory Contention**, **CPU Ready**, and **CPU Co-Stop** and see if there is any color other than green.

- Green indicates that almost 100% of the VMs have received the CPU and memory that was requested. The threshold is set such that if the 10% of the VM population does not receive the requested resources, then the heat map turns red.
- Red indicates an early warning. Stringent thresholds are used to enable proactive attention and remediation operations. The heat map can turn red because of the high standard that is applied even when there is no complaint from the VM owner yet.
- The light gray indicates that there is no VM running on the host and the metric is not computing.

View if there is any unbalance.

- There are two types of unbalance, cluster unbalance, and resource type unbalance.

- The ESXi hosts are grouped by the cluster, so that the unbalance within a cluster can be easily viewed. Cluster unbalance is a real possibility and it is best monitored and not just assumed.
- If the three heat maps are different, then there is a resource unbalance. For example, if the memory contention is mostly red, but the two CPU heat maps are green, it means you have an unbalance between memory and CPU.
- If a single ESXi host displays different color across the three heat maps, it indicates that there is an unbalance between the CPU and memory resources in the host.

For NOC Operator, drill-down by selecting one of the VMs on the heat map.

- The **Trends of Selected ESXi Host** widget will automatically display the performance counters. To hide any metric, click the name in the legend.

As part of the deployment, configure auto-rotate among the NOC dashboards. If you want to view one dashboard, then you can remove the vRealize Operations Manager menu by using the URL sharing feature. This makes the overall user interface presentable and allows you to focus on the dashboard.

Points to Note

- You can add Disk Latency if you have the screen real estate. Use the counter 'Percentage of Consumers facing Disk Latency (%)'. It is a part of a datastore object, not a cluster, as a VM in a cluster can have disks across multiple datastores. Organize this storage performance by data center and not by the cluster.

Live! Cluster Utilization Dashboard

The **Live! Cluster Utilization** dashboard complements the **Cluster Performance** dashboard. Use this dashboard to view the clusters that are working excessively and are close to their physical limit. This dashboard displays ESXi hosts that have CPU or memory saturation that can lead to performance issues for the VMs running on the host.

Design Considerations

This dashboard is designed to complement the **Live! Cluster Performance** dashboard and it shares design considerations.

How to Use the Dashboard

As this dashboard has an identical design with the **Live! Cluster Performance** dashboard, it has the same usage procedure. Unlike the heat maps in the **Live! Cluster Performance** dashboard, the three heat maps in this dashboard have a different scale, reflecting the different nature of the counters.

Logically, memory is a form of storage. It acts as a cache to disk as it is much faster. A high utilization is better, as it indicates that more data is being cached. The ideal situation is when ESXi host Consumed metric is red but ESXi host Ballooned metric is green. When Ballooned is red and Consumed is gray, it means that there was high pressure in the past but it is not there anymore. The reason the ballooned stays red is because the ballooned pages were never requested back.

The ballooned memory counter was selected over the swapped or compressed memory counters as it is a better leading indicator. Since all three can co-exist at the same time, they are displayed in the line chart. Ballooned is displayed in absolute amount and not as a percentage, because the higher the size, higher are the chances for it to impact a VM. If you feel using percentage is easier for your operations, create a super metric to translate the value.

The heat map displays Wastage by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

Analyze if the ESXi host is contributing. A light gray box indicates that the host is a part of the cluster but there is no utilization. It is possible for the host to be in the maintenance mode or is powered off.

Points to Note

- ESXi host chooses to swap over compression if the compression ratio is less than 4x.
- If the ESXi host's physical NIC is saturated in your environment, then you can add a Network Throughput heat map.

Live! Heavy Hitters Dashboard

The **Live! Heavy Hitters** dashboard helps you analyze the misuse of the shared infrastructure. This dashboard displays details of VMs misusing shared infrastructure and if that has caused performance problems to the other VMs. The shared infrastructure includes risks. The cause for excessive load might be attacks, for example, the denial of service, process runaway, or a mass activation of agents. The most demanding VM is the largest. If a handful of VMs is dominating the shared infrastructure, their collective size is displayed on the dashboard.

Design Considerations

See the [Performance Dashboards](#) page for common design considerations among all the dashboards for performance management.

In a shared environment, it is possible to have a victim-villain problem. In the heat map, the villain VM is the one with the largest box size, while the victim VM is the one with the red box. If a handful of VMs is dominating the shared infrastructure, their collective size will be highly visible on the dashboard.

How to Use the Dashboard

- The heat maps, Disk IOPS, Disk Throughput, Network Throughput, and CPU Demand displays the four different loads that can be excessive. The heat maps display the relative value and not the absolute value. A VM does not generate a high load in the absolute term just because it has a large configuration.
- Each heat map has its color threshold, reflecting the nature of the contention metrics used in each of them.

- For NOC Operator, drill-down by selecting one of the VMs on the heat map. All the four line charts are automatically displayed, enabling you to get a complete picture of the selected VM.

Points to Note

- Memory is not displayed as it is a form of storage. The memory counters are space utilization and not speed. Think of disk space instead of IOPS. It can cause a capacity problem on the shared ESXi host, but not performance problems to other VMs.
- In a large environment, it might be difficult to view a small victim VM. Consider having multiple dashboards and use them interchangeably.

Software Defined Wide Area Network Dashboard

The Software-Defined Wide Area Network (SD-WAN) dashboard allows you to configure and monitor the services related to VeloCloud and SD-WAN using vRealize Operations Manager. Using the SD-WAN dashboard, you can also collect the metrics for VeloCloud Orchestrator and VeloCloud Gateway.

By default the SD-WAN dashboards are disabled, if you want to know how to enable them, see [Manage Dashboards](#). The following services are discovered using the VeloCloud Orchestrator:

- Java Application
- VeloCloud Orchestrator
- Nginx
- ClickHouse
- MySQL
- Redis
- Network Time Protocol

The following services are discovered using VeloCloud Gateway:

- Network Time Protocol
- VeloCloud Gateway

Troubleshoot SD-WAN Dashboard

You can use the widgets in Troubleshoot SD-WAN dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN.

You can use the dashboard widget in several ways:

- **Troubleshoot Virtual Machine (VM):** Use this widget to navigate to a specific VM and troubleshoot the issues.
- **Troubleshoot Orchestrator:** Use this widget to navigate to a specific orchestrator and troubleshoot the issues.

- **Troubleshoot Gateway:** Use this widget to navigate to a specific gateway and troubleshoot the issues.
- **Troubleshoot Application:** Use this widget to navigate to a specific application and troubleshoot the issues.
- **Relationship:** Use this widget to view the services and operating system associated with the VeloCloud Orchestrator.
- **Top Alerts:** Use this widget to view the top alerts associated with the SD-WAN.

Troubleshoot SD-WAN Gateway Dashboard

You can use the widgets in Troubleshoot SD-WAN Gateway dashboard to monitor and troubleshoot all the services and applications associated with the SD-WAN gateway.

You can use the dashboard widget in several ways:

- **Active Alerts on the Gateway:** Use this widget to view the active alerts for the gateway.
- **Health of Gateway Applications:** Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System:** Use this widget to examine the operating system status.
- **Gateway Summary Status:** Use this widget to view the summary information for the gateway.
- **Gateway Process Status:** Use this widget to view the process information for the gateway.
- **Gateway Resource Metrics:** Use this widget to view the resource metrics associated with the gateway.
- **Parent Host:** Use this widget to view the parent host information.
- **Parent Cluster:** Use this widget to view the parent cluster information.

Troubleshoot SD-WAN Orchestrator Dashboard

You can use the widgets in Troubleshoot SD-WAN Orchestrator dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN Orchestrator.

You can use the dashboard widget in several ways:

- **Active Alerts on the Orchestrator:** Use this widget to view the active alerts for the Orchestrator.
- **Health of Orchestrator Applications:** Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System:** Use this widget to examine the operating system status.
- **Examining MySQL:** Use this widget to examine the MySQL application.
- **Orchestrator Service Status:** Use this widget to view the service status of the Orchestrator.
- **Redis Status:** Use this widget to view the status of the Redis application.

- **API Check Status:** Use this widget to check the API status.
- **Nginx Status:** Use this widget to check the Nginx status.
- **Parent Host:** Use this widget to view the parent host information.
- **Parent Cluster:** Use this widget to view the parent cluster information.

vRealize Automation 8.x Dashboards

With the vRealize Automation 8.x dashboards, you can monitor and troubleshoot objects in your cloud infrastructure.

The following vRealize Automation 8.x dashboards are added to the predefined vRealize Operations Manager dashboards:

- Cloud Automation Environment Overview
- Cloud Automation Project Cost Overview
- Cloud Automation Resource Consumption Overview
- Cloud Automation Top-N Dashboard

Cloud Automation Environment Overview

You can use the widgets in the Cloud Automation Environment Overview dashboard to view the environment details for the vCenter Cloud Zone objects. You can use the Cloud Automation Environment Overview dashboard to view the projects, deployments associated with the vCenter Cloud accounts.

You can use the dashboard widgets in several ways.

- **vCenter Cloud Zone List:** Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the cloud zone objects present in your environment.
- **Project List:** Use this widget to view the total blueprints, cloud zones, deployments, virtual machines, health, risk, efficiency details in your environment.
- **Top Alerts:** Use this widget to view the top alerts in your environment.
- **VM List:** Use this widget to view all the VM details in your environment.
- **Blueprint List:** Use this widget to view the blueprint objects in your environment.
- **Deployment List:** Use this widget to view the blueprint objects deployed in your environment.

Cloud Automation Project Cost Overview

You can use the widgets in the Cloud Automation Project Cost Overview dashboard to view the project cost associated with cloud zone objects present in your environment.

You can use the dashboard widgets in several ways.

- **Project Cost:** Use this widget to view the project wise cost for compute, storage, and additional resources associated with your cloud environment.
- **Total Cost Over Time:** Use this widget to view the cost of individual projects on a day to day basis.
- **Deployment Cost by Selected Project:** Use this widget to view the deployment cost for the selected project in your cloud environment.

Cloud Automation Resource Consumption Overview

You can use the widgets in the Cloud Automation Resource Consumption Overview dashboard to view the resources consumed by vRealize Automation 8.x on Cloud Accounts.

You can use the Cloud Automation Resource Consumption Overview dashboard widgets in several ways.

- **Cloud Account:** Use this widget to view all the attributes related to the cloud account.
- **Cloud Zone:** Use this widget to view all the attributes related to the cloud zones.
- **Project:** Use this widget to view all the project details associated with your cloud account.
- **Cluster List:** Use this widget to view all the details associated with the clusters in your account.
- **Cluster Utilization:** Use this widget to view the cluster utilization details for the cloud accounts.
- **Deployment Heat Map by Project:** Use this widget to view the heat map for each deployed project in your cloud environment.
- **Cloud Zone Capacity:** Use this widget to view the memory and storage capacity that is allocated, reserved, and free for each cloud zone object.
- **Cloud Zone Memory Trend:** Use this widget to view and analyze a seven-day trend for the memory allocated, reserved, and free for the cloud zone.
- **Cloud Zone Storage Trend:** Use this widget to view and analyze a seven-day trend for the storage allocated, reserved, and free for the cloud zone.

Cloud Automation Top-N Dashboard

You can use the widgets in the Cloud Automation Top-N dashboard to view the projects with most critical alerts, to view the blueprint with most deployments, and to view the deployments with the highest cost.

You can use the dashboard widgets in several ways.

- **Project with Most Critical Alerts:** Use this widget to view the projects which has most critical alerts.
- **Top Alerts:** Use this widget to view the top alerts for the projects in your cloud account.

- **Blueprints with Most Deployment:** Use this widget to view the blueprint which has maximum deployments for the cloud account.
- **Relationship:** Use this widget to analyze the relationship between blueprints and deployments, and deployment and cost.
- **Deployment with Highest Cost:** Use this widget to identify the most expensive deployment associated with your cloud account.

Service Discovery Dashboards

Using the service discovery dashboards, you can determine the inter-dependencies of virtual machines and the dependencies of each service in the respective virtual machines.

The following service discovery dashboards are added to the predefined vRealize Operations Manager dashboards:

- Service Distribution
- Service Relationships
- Service Visibility
- Virtual Machine Relationships

Service Distribution Dashboard

You can use the dashboard to view the distribution of different services in the selected data center, cluster, or a host system. You can also view known and unknown services including the category and distribution percentage across a vSphere resource.

You can use the dashboard widgets in several ways:

- **Inventory Item:** Use this widget to view a hierarchical representation of objects in the form of badges.
- **Known Services Distribution:** Use this widget to view different services discovered from a selected object.
- **Service Categories:** Use this widget to view the service categories that are discovered by selecting an object from the resource widget.
- **User Defined Services Distribution:** Use this widget to view a list of user-defined services.

Service Relationships Dashboard

You can use the dashboard to view properties of the service such as the install path, the ports used, and the version. You can also view the relationship between the services that run on other VMs.

You can use the dashboard widgets in several ways:

- **List of Services Discovered:** Use this widget to view the services that have been discovered.

- **Connections from the Selected Services:** Use this widget to view the relationship between the services and the other services running on the VMs.
- **Properties of the Selected Service:** Use this widget to view the properties of the selected services.

Service Visibility Dashboard

You can use the dashboard to view a list of VMs without service visibility and VMs with user-defined services after you select a vSphere object.

You can use the dashboard widgets in several ways:

- **Inventory Tree:** Use this widget to view a hierarchical representation of objects in the form of badges.
- **Virtual Machines without Service Visibility:** Use this widget to view information about services where discovery has failed.
- **Virtual Machines with User-Defined Services:** Use this widget to view a list of VMs where the user has defined such services.

Virtual Machine Relationships Dashboard

You can use the dashboard to view a list of VMs with service discovery details such as, status, method, incoming/outgoing connections, and protection groups. When you select a VM, the dashboard displays a list of discovered services on the VM, the relationships of the VMs with other VMs based on the relationships of the discovered service.

You can use the dashboard widgets in several ways:

- **List of virtual machines:** Use this widget to view all the VMs discovered by the vCenter Server.
- **Node relationship of the selected VM:** Use this widget to view the relationship between the objects.
- **List of Services running in the selected VM:** Use this widget to view all the properties of the selected VM.
- **Connections of Virtual Machines:** Use this widget to view the relationship between one or more VMs.

Inventory Dashboards

The three vSphere Inventory dashboards and workload management inventory dashboards cater to the compute, network, and storage aspects of your SDDC. Using these dashboards, you can navigate through the environment and view your inventory and their key metrics at a glance. The Network and Storage dashboards can be shared with the network and storage teams respectively, giving them the necessary visibility, and increasing the collaboration between teams.

vSphere inventory dashboards

The vSphere inventory dashboards are built specifically for each role, but they share a common design. They have a similar layout and are used in the same manner. This makes learning easier, especially in smaller environments where the same team manages the full environment.

These dashboards help you answer several key questions:

- What is the topology of your vSphere compute inventory?
- What is the topology of your vSphere storage inventory?
- What is the topology of your vSphere network inventory?

Workload Management Inventory Dashboard

This is a unified dashboard for the new workload management objects. It shows the relationships and KPIs for the workload management objects. For example, you can see the topology view from the Tanzu Kubernetes clusters to the physical infrastructure.

vSphere Compute Inventory Dashboard

You can use the vSphere Compute Inventory Dashboard to browse through the topology of your vSphere compute inventory which includes information related to vSphere world, vCenter Server, data center, clusters, hosts, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the clusters, ESXi hosts, and virtual machines associated with the object.

You can use the dashboard widgets in several ways.

- **Properties:** View the properties related to an object in the environment.
- **Metrics:** View the metrics related to the object.
- **Clusters:** View the cluster functionality.
- **ESXi Hosts:** View the data related to the hosts.
- **Virtual Machines:** View VMs that belong to the object.

vSphere Network Inventory Dashboard

The vSphere Network Inventory Dashboard allows you to browse through the topology of your vSphere network inventory which includes information related to vSphere world, vCenter Server, data center, distributed vSwitches, distributed port groups, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the distributed vSwitches, distributed port groups, virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties:** View the properties related to the object in the environment.

- **Metrics:** View the metrics of the object.
- **Distributed vSwitches:** View details related to the distributed vSwitches.
- **Distributed Port Groups:** View data relevant to distributed port groups.
- **Virtual Machines:** View VMs that belong to the object.

vSphere Storage Inventory Dashboard

The vSphere Storage Inventory dashboard allows you to browse through the topology of your vSphere storage inventory which includes information related to vSphere world, vCenter Server, data center, datastore clusters, datastores, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the datastore clusters, datastores, and virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties:** View the properties related to the object in the environment.
- **Metrics:** View the metrics of the object.
- **Datastore Clusters:** View the datastore cluster functionality.
- **Datastores:** View the datastore functionality.
- **Virtual Machines:** View VMs that belong to the object.

Workload Management Inventory Dashboard

The Workload Management Inventory dashboards curates the Kubernetes inventory across all the Workload Management enabled vSphere environments and displays it here. This includes an end to end topology map showcasing the health of all the objects along with upstream and downstream dependencies. Upon clicking any object in the relationship tree, the related inventory of Supervisor Clusters, Namespaces, Pods, Developer Managed VMs and Tanzu Kubernetes clusters can be viewed and exported from this dashboard.

You can select an object type to view the properties and key metrics related to it.

You can use the dashboard widgets in several ways.

- **Environment Summary:** Provides a summary of the supervisor cluster and the child objects.
- **Relationships:** An interactive canvas where you can view the relationship between the different objects in the workload management inventory.
- **Properties:** View the properties related to the object in the environment.
- **Metrics:** View the metrics of the object.
- **Supervisor Clusters:** View the supervisor cluster functionality.
- **Tanzu Kubernetes cluster:** View the Tanzu Kubernetes cluster functionality.
- **Virtual Machines:** View VMs that belong to the object.

- **vSphere Pods:** View information about vSphere Pods.

Microsoft Azure Dashboards

Use dashboards to monitor and troubleshoot Microsoft Azure issues in vRealize Operations Manager.

To access the dashboards, click **Dashboards** on the menu and click the dashboard names that start with Azure.

The following dashboards are available:

Dashboard Name	Purpose
Availability	View the availability of each Microsoft Azure service. Available services are green. Unavailable services are red and will be removed.
Inventory	<p>View the adapter instance count in each resource group. Select a resource group to see a sparkline chart and the metrics for all the resources in the group.</p> <p>Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the inventory for the database.</p> <p>Note Metrics that are not collected or created are grayed out.</p>
Optimization	View whether you are effectively using Microsoft Azure services. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
Virtual Machine	Select a virtual machine to view its scoreboard, property list, object relationship with resource group, and CPU usage and forecasting. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
SQL Database	Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the scoreboard, object relationship, and CPU usage for the database. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.
Load Balancer	Select a load balancer to view its scoreboard, object relationship, and data path availability. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart.

AWS Dashboards

Dashboards provide the user interface you use to monitor and troubleshoot Amazon Web Services problems in vRealize Operations Manager.

You can access the dashboards by selecting **Dashboards**, and then selecting **AWS**.

Table 5-1. AWS Dashboards

Dashboard Name	Purpose
AWS Alerts	The Alerts dashboard reports system-generated performance information for Amazon Web Services. In vRealize Operations Manager 5.8 and later, the dashboard also displays alerts received from Amazon Web Services Cloudwatch.
AWS ASG Utilization	Use the Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric. ASG metrics are not collected by default. You must enable them when creating the group. This applies only to the metrics belonging directly to the auto scale group, for example GroupDesiredCapacity. It does not apply to the aggregate instance metrics for the ASG, for example Instance Aggregate CPU Utilization.
AWS Disk Space	Use the Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default. For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html
AWS Instance Heatmap	Use the Instance Heatmap to monitor CPU/Disk/Network metric elements and identify instances that perform poorly.
AWS Instance Utilization	Use to identify which EC2 instances have high use across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by making adjustments to EC2 instances.

Table 5-1. AWS Dashboards (continued)

Dashboard Name	Purpose
AWS Troubleshooting	<p>This dashboard is most helpful when someone calls in with a problem and you know which device they are using. You can search for that type of device or the specific device, if you know the name.</p> <p>When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas. You can use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another.</p>
AWS Volume Performance	Use the Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations.
AWS Availability	Use this dashboard to view the availability of each AWS service.
AWS Inventory	Use this dashboard to view the count of each AWS service instance in each region.
AWS Optimization	Use this dashboard to view if you are effectively using AWS services.

Table 5-2. AWS - All Other Dashboards

Dashboard Name	Purpose
<p>AWS Services</p> <ul style="list-style-type: none"> ■ CloudFormation Stacks ■ Compute: EC2 ■ Compute: Elastic Containers ■ Compute: Lambda Functions ■ Database: Dynamo ■ Database: ElastiCache ■ Database: RDS ■ Database: Redshift ■ Desktop: Workspaces ■ Network: Load Balancers ■ Network: VPS ■ Simple Queue Services ■ Storage 	Select AWS Services and then select a dashboard to view a specific service-related information.

AWS Instance Utilization Dashboard

Use the AWS Instance Utilization dashboard to identify which EC2 instances have a high usage across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by adjusting the EC2 instances.

For example, you might determine that you need to resize the EC2 instance to make it larger or smaller.

You most often use this dashboard to troubleshoot issues with the listed metrics based on a support request from a user.

You can also identify which EC2 instances have been running for the longest and shortest amount of time. Then, you can use that information to determine whether EC2 instances can be decommissioned, or discover instances that have been added and need to be tracked in inventory.

Memory metrics require that you implement an add-on for each EC2 instance. These add-ons cost extra, and are not included by default.

AWS Auto Scaling Group Dashboard

Use the AWS Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric.

AWS Troubleshooting Dashboard

When a user calls in with a problem and you know the name of the device they are using, can search for that type of device or the specific device and use the AWS Troubleshooting dashboard to get an overview of the system functionality.

When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas.

Use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another.

There is a suggested flow to using the widgets in this dashboard.

- 1 Start with only the AWS Object widget open, and find the item you want to inspect.
- 2 Select the item, then expand the AWS Relationship widget to view the item's status.
- 3 Select one or all the related objects, then view the Ordered Symptoms, Interesting Metrics, and Mash-up.

- 4 Optionally, drag widgets into a new configuration if it makes it easier for you to compare information that is meaningful to you.
- 5 Examine the list of ordered symptoms and determine which of these events, in the given order might cause the problem to occur.

AWS Instance Heatmap Dashboard

Use the AWS Instance Heatmap dashboard to monitor CPU/Disk/Network metric elements and identify instances that perform poorly.

You can use the Troubleshooting dashboard to find more detail, and research the root cause of issues. Then you can view the specific object instance to identify faulty processes and take a corrective action.

AWS Volume Performance Dashboard

Use the AWS Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations. When you identify the EC2 instance that generates the load, use the Troubleshooting dashboard to investigate further.

AWS Disk Space Dashboard

Use the AWS Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default.

For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html>.

AWS Alerts Dashboard

The AWS Alerts dashboard reports system-generated performance information for Amazon Web Services. In vRealize Operations Manager 6.6 and later, the dashboard also displays alerts received from Amazon Web Services Cloud watch.

Dashboards in VMware Cloud on AWS

The **VMware Cloud on AWS** dashboards allow you to track the capacity, cost, and inventory overviews of the SDDCs. You can also track the virtual machines monitoring and the utilization and performance of these SDDCs.

VMC Capacity Dashboard

Use the **VMC Capacity** dashboard to view the capacity overview of each VMware Cloud on AWS SDDC. You can view the capacity of Clusters, Hosts, VMs, Datastores, and Disk groups.

Table 5-3. Widgets in VMC Capacity Dashboard

Widget	Description
VMC SDDC by Capacity Remaining %	Displays the SDDCs as cards that show the remaining capacity percentage.
VMC SDDC by Time Remaining %	Displays the SDDCs as cards that show the remaining time percentage.
VMC SDDC by Virtual Machine Remaining (based on avg VM profile)	Displays the SDDCs as cards that show the remaining number of virtual machines.

When you select one of the SDDC cards, the details of that SDDC are automatically populated in the widgets after the VMC SDDC by Virtual Machine Remaining (based on avg VM profile) widget.

Note The key kpis are color-coded to help in identifying capacity bottlenecks.

VMC Cost Overview Dashboard

Use the **VMC Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

Table 5-4. Widgets in VMware Cloud on AWS Dashboard

Widget	Description
Organization Cost Overview	Displays a list of organizations with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD).
Outstanding Expense Trend	Displays the outstanding expense trend of the organization selected in the Organization Cost Overview widget.
Total Expense Trend (Monthly)	Displays the total monthly expenses trend of the organization selected in the Organization Cost Overview widget.
Commit Expense Trend (Monthly)	Displays the committed monthly expense trend of the organization selected in the Organization Cost Overview widget.
On-Demand Expense Trend (Monthly)	Displays the on-demand monthly expenses trend of the organization selected in the Organization Cost Overview widget.
Purchase History	Displays the bill line items/purchases from the available bills.
Currency Information	Represents the metrics currency unit set in this management pack account.

Note The YTD metric is an aggregation from the beginning of the calendar year, until the last available bills.

VMC Inventory Dashboard

Use the **VMC Inventory** dashboard to view the inventory overview of all the SDDCs configured in VMware Cloud on AWS.

Widgets in VMC Inventory Dashboard

VMC SDDCs: displays the SDDCs as cards that show the number of virtual machines running in the SDDC. The SDDC card also shows a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that SDDC, the SDDC card indicates this by changing colors.

When you select one of the SDDC cards, the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs with key configuration details of that SDDC are populated in the widgets after the VMC SDDCs widget.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

VMC Management VM Monitoring Dashboard

Use the **VMC Management VM Monitoring** dashboard to monitor the utilization and performance of the key management VMs running in your SDDC. This dashboard ensures that the management components (such as vCenter and NSX) are not facing any resource bottlenecks from the CPU, memory, network, and storage perspectives.

Table 5-5. Widgets in VMC Management VM Monitoring Dashboard

Widget	Description
CPU Usage & Performance	Displays the list of all the management components in each SDDC with key CPU utilization and performance KPIs. Select a management VM to see the usage and performance trends of all the CPU cores.
Memory Usage & Performance	Displays the list of all the management components in each SDDC with key Memory utilization and performance KPIs. Select a management VM to see the memory usage and performance trends.
Network Usage & Performance	Displays the list of all the management components in each SDDC with key Network utilization and performance KPIs. Select a management VM to see the memory usage and performance trends.
Storage Usage & Performance	Displays the list of all the management components in each SDDC with key storage utilization and performance KPIs. Select a management VM to see the network usage and performance trends.

VMC Utilization and Performance Dashboard

Use the **VMC Utilization and Performance** dashboard to view the utilization and performance overview of each SDDC based on heavy hitter VMs and impacted VMs over the last 30 days. This dashboard helps you in finding the VMs in your environment that are negatively impacting the capacity or performance from a CPU, memory, storage, or network perspective.

Widgets in VMC Utilization and Performance Dashboard

List of VMC SDDCs: displays the list of all the SDDCs with aggregate CPU, memory, and storage utilization with 95th percentile and maximum values over the last 30 days.

When you select one of the SDDC from the List of VMC SDDCs widget, you can see the list of top VMs that are consuming compute, network & storage resources in that SDDC. The widgets after that show the compute (CPU & memory) utilization and performance analysis, network, storage, and utilization and performance analysis.

Each section in the dashboard is based on the last 30 days data with 95th percentile transformation which is configurable to Max, Average, Current, Standard Deviation, or other mathematical transformations.

VMC Configuration Maximums Dashboard

Use the **VMC Configuration Maximums** dashboard to view the VMC limits and your consumption against those limits. This dashboard displays alerts for configuration maximum, and details of organization, SDDC, vSAN, and cluster maximums.

Table 5-6. Widgets in VMC Configuration Maximums Dashboard

Widget	Description
Select an Environment	Select an environment for which you want to view the alerts and other details. Once you select an environment, the details of that environment are automatically populated in the widgets below.
VMC Configuration Maximums Alerts	Displays the list of alerts for the selected environment.
Number of SDDCs	Displays the number of SDDCs for the organization maximums, the provisioned, and the soft limit used.
Number of Hosts	Displays the number of hosts for the organization maximums, the provisioned, and the soft limit used.
Public IP Addresses (Elastic IPs)	Displays the public IP addresses for the organization maximums, the provisioned, and the soft limit used.
Maximum Clusters	Displays the maximum clusters for the SDDC maximums, the provisioned, and the hard and soft limit used.
Maximum Hosts	Displays the maximum hosts for the SDDC maximums, the provisioned, and the limit used.
Maximums VMs	Displays the maximum VMs for the SDDC maximums, the provisioned, and the limit used.
Linked VPCs	Displays the linked VPCs for the SDDC maximums, the provisioned, and the limit used.
Clusters with No SLA	Displays the maximum number of clusters and the number of provisioned clusters with no SLA per SDDC. An empty list means no clusters have been identified with no SLA.
Clusters with Limited SLA	Displays the maximum number of clusters and the number of provisioned clusters with limited SLA per SDDC. An empty list means no clusters have been identified with limited SLA.
Max hosts per Cluster (including stretched clusters)	Displays the maximum hosts per cluster including the stretched clusters, the provisioned, and the limit used.
Datastore Utilization	Displays the datastore utilization for vSAN maximums, the used space, utilization limit, and the remediation needed.

Table 5-6. Widgets in VMC Configuration Maximums Dashboard (continued)

Widget	Description
VMs per Host Limit Used	Displays the maximum number VMs that can be deployed per host, VMs that are provisioned per host, and the percentage of limit used.
VMs per Host Limit Used of Selected Host	Displays the VMs that are used per host limit for a selected host.

Dashboards in NSX-T Management Pack

The **NSX-T Main** dashboard provides an overview of the network objects. It displays the topology of a selected object, how it connects to the elements in the network, and a view of related alerts.

Table 5-7. Widgets in NSX-T Main Dashboard

Widget	Description
NSX-T Instances	Displays the list of environments that are being monitored. When you select an environment in this widget, the other widgets in the NSX-T Main dashboard display data for the selected adapter.
Environment Overview	Displays a top-level view of the selected environment and following key components. <ul style="list-style-type: none"> ■ NSX-T Manager ■ Controller Node ■ Logical Router ■ Logical Switch ■ Load Balancer Virtual Server ■ Transport Zone
Top Alerts	Displays all the open alerts for the selected object in the Environment Overview widget.
Topology Graph	Displays the topology of the selected object in the Environment Overview widget.

NSX-T Configmax Metrics

The **NSX-T Configmax Metrics** dashboard provides an overview of all the configuration maximum metrics in all the NSX-T instances.

Table 5-8. Widgets in NSX-T Configmax Metrics Dashboard

Widget	Description
Select an adapter instance	Displays the list of all the NSX-T and NSX-T on VMC instances. When you select an instance in this widget, the other widgets in the NSX-T Configmax Metrics dashboard display data for the selected instance.
Relationship view	Displays the objects hierarchy for the instance selected in the Select an adapter instance widget. Only the objects with configuration maximum metrics are shown in the relationship view.

Table 5-8. Widgets in NSX-T Configmax Metrics Dashboard (continued)

Widget	Description
Select object from relationship view for the configmax metric	Displays all the configmax metrics for the selected object in the Relationship View widget.
Trend View	Displays all the MGW, CGW, and Distributed firewall section rule trends of the instance selected in the Select an adapter instance widget. Note The Trend View widget loads the trends only for the firewall sections object on VMware Cloud on AWS instances.

Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager

6

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.

When your customers experience performance problems and call you to resolve the problem, the data that vRealize Operations Manager collects and processes is presented to you in graphical forms. You can then compare and contrast objects, understand the relationship between objects, and determine the root cause of problems.

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 change objects in the target system, for example, the VMware vCenter Server® system, from vRealize Operations Manager .

This chapter includes the following topics:

- [Enhanced Search Capability](#)
- [What to Do When...](#)
- [Troubleshooting Workbench Home Page](#)
- [Monitoring and Responding to Alerts](#)
- [Monitoring and Responding to Problems](#)
- [Running Actions from vRealize Operations Manager](#)
- [Viewing Your Inventory](#)

Enhanced Search Capability

The search function on the upper right supports locating named objects, dashboards, alerts, and so on, in the system. The search function attempts to match or partially match any string you enter; additional capabilities enable you to go swiftly to the item you want. The system presents the item in the Edit context.

Where you Find Search

The search function appears on all the pages of the vRealize Operations Manager in the top menu. Click the magnifying glass icon to open the search bar. Optionally, you can press the Ctrl, Shift and Spacebar keys on your keyboard to open the search bar.

How Search Works

You start your search by typing in the search bar. vRealize Operations Manager displays matching objects types and objects.

The search function supports several common categories you can employ to find the item you seek quickly, as follows:

- Dashboard
- Object
- Supermetric
- Alert definition
- Symptom definition
- View
- Report
- Notification
- I.P. Address

What this means is that in addition to entering a traditional search phrase, for example, a simple string - "VM" - you can also enter one of the listed categories followed by a string or a name. You can then search for objects within the category. For the Object, View and Dashboard categories, the system displays the object in view mode.

If you want quickly to locate a specific dashboard, for example, start typing "dash..." into the search field. The system offers the search term Dashboards. Select the term using the cursor and then enter the dashboard name or part of the name and press Enter. The system finds the dashboard you want, with editing functions available.

Similarly, you can type "alert" or simply "a" in the search field and the system offers Alert Definition. Select the term and enter part of an alert message, for example, "unbalanced." The system returns the alert, "Cluster has an unbalanced workload," presented in the Alert Definition Workspace where you can edit it.

Note You can type virtual machine in the search bar to list all the virtual machines associated with the host.

What to Do When...

As a virtual infrastructure administrator, network operations center engineer, or other IT professional, use vRealize Operations Manager to monitor objects in your environment. Using vRealize Operations Manager, you can ensure that your customers experience the best possible service, 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 tech support reporting that a virtual machine, VPSALES4632, is running slowly. The VP is working on sales reports for an upcoming meeting and is running behind schedule because of the slow performance of the 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 act to resolve the problems.

User Scenario: A User Calls with a Problem

The vice president of sales telephones tech support reporting that a virtual machine, VPSALES4632, is running slowly. The VP is working on sales reports for an upcoming meeting and is running behind schedule because of the slow performance of the virtual machine.

As an operations engineer, you reviewed the morning alerts and did not see problems with that 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 Troubleshooting to Investigate a Reported Problem

To troubleshoot problems with the VPSALES4632 virtual machine, consider evaluating symptoms, examining time line information and events, and creating metric charts to find the root cause.

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 on the vRealize Operations Manager title bar, enter the name of the virtual machine.

The **Search** text box displays all the objects that contain the string you enter in the text box. If your customer knows that the virtual machine name contains SALES, enter 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.
- If you cannot locate the cause of the reported problem among the alerts, begin more in-depth troubleshooting. See [Use Troubleshooting to Investigate a Reported Problem](#).

Use Troubleshooting to Investigate a Reported Problem

To troubleshoot problems with the VPSALES4632 virtual machine, consider evaluating symptoms, examining time line information and events, and creating metric charts to find the root cause.

If a review of the alerts did not help you identify the cause of the problem reported for the virtual machine, use the following tabs: **Alert > Symptoms**, **Event > Timeline**, and **All Metrics** to troubleshoot the virtual machine history and current state.

.

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 use?

- 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.

It appears 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 that might help 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. Metrics might direct you toward the cause of the reported problem.
 - a Manipulate the **Date Controls** to identify the approximate time when your customer reported the problem.
 - b Use the Filters to filter on event criticality and status. Select Symptoms if you want to include the filters 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 problems are the likely cause of the problem.

- 5 If the problem relates to CPU or memory use, click **All Metrics** and create metric charts to identify whether it is CPU, memory, or both.
 - a If the host is still the focus, begin 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.

Comparing the four charts shows normal CPU use on both the host and the virtual machine, and normal memory use on the virtual machine. However, memory use on the host is consistently elevated three days before the reported problem on VPSALES4632.

Results

The host memory is consistently elevated, which impacts virtual machine response time. The number of running virtual machines is well within the supported number. The cause might be many intensive process applications on the virtual machines. 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, use vRealize Operations Manager to power off virtual machines on the host so that you can improve performance in the running virtual machines. See [Run Actions from Toolbars in vRealize Operations Manager](#) .
- If you want to use the combination of charts that you created on the **All Metrics** tab 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. You receive emails when an alert is generated for your monitored objects. In addition to alerting you to problems in your environment, alerts can 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. It also assumes 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 messages when an alert is generated so that you can respond quickly.

Prerequisites

- Verify that outbound alerts are configured for standard email alerts. See *Add a Standard Email Plug-In for vRealize Operations Manager Outbound Alerts* topic in *vRealize Operations Cloud Configuration Guide*.
- Verify that outbound alerts are configured for standard email alerts. See *vRealize Operations Cloud Configuration Guide*.
- 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* topic in *vRealize Operations Cloud Configuration Guide*.

Procedure

1 Respond to an Alert in Your Email

As a network operations engineer, you receive an email message from vRealize Operations Manager about a datastore 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 Datastore

Because you need more information about the datastore before you decide on the best response, you examine the **Symptoms** tab to see other triggered symptoms for the datastore.

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 Datastore 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 solve 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 about a datastore 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 from 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 use reaching limit  datastore1      | b0885859-
e0c5-4126-8eba-6a21c895felb      | 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* topic in *vRealize Operations Cloud Configuration Guide*.
- Verify that outbound alerts are configured for standard email alerts. See *vRealize Operations Cloud Configuration Guide*.
- 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 an Email Alert Notification* topic in *vRealize Operations Cloud Configuration Guide*.
- 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 *vRealize Operations Cloud Configuration Guide*.

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 problem. If implemented, do the prioritized recommendations resolve the problem?
What is Causing the Problem?	Which symptoms were triggered? Which were not triggered? What effect 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 might 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 datastore. See [Evaluate Other Triggered Symptoms for the Affected Datastore](#).

Evaluate Other Triggered Symptoms for the Affected Datastore

Because you need more information about the datastore before you decide on the best response, you examine the **Symptoms** tab to see other triggered symptoms for the datastore.

If other symptoms are triggered for the object besides the symptom included in the alert, evaluate them as well. Determine what the symptoms reflect about the state of the object 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.
- 2 In the **Alert Details** tab, see the information under **Symptoms**. Click the object which is displaying the symptoms.
- 3 The object opens under **Environment**. Click **Alerts > Symptoms**. The symptoms tab includes all the symptoms triggered for the current object.

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 that might 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 can 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 can 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 Datastore Alert](#).

Construct Metric Charts to Investigate the Cause of the Datastore 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 can resolve the problem that the alert identifies.

Prerequisites

View the topological map for the datastore to determine if related objects are contributing to the alert or if triggering symptoms indicate that the datastore 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, the Capacity | Used Space (%) chart might show an increase in used space, without the Disk Space | Capacity (GB) increasing or the Summary | Number of Capacity Consumers increasing. Then adding capacity can be 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 four charts.

For example, the Summary | Total Number of VMs chart might show that the number of virtual machines did not increase enough to affect the datastore negatively. That result might make moving some of the virtual machines seem the best 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, say the amount of Disk Space | Snapshot Space (GB) increases. At the same time, the Disk Space Reclaimable | Snapshot Space | Waste Value (GB) indicates an area where space can be reclaimed. Then deleting unused snapshots positively affects the datastore disk space problem and resolves the alert.

- 5 If this datastore is a problematic one that you must continue to monitor, 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 recommendation to Delete unused snapshots of Virtual Machines 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 solve 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 task deletes the snapshots. The status is `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 act 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 emailed, 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 analyze the state of the objects in your environment, examine current problems, investigate solutions, and act to resolve the problems.

This scenario shows you how to evaluate the problems that occur on your objects, and how to resolve problems.

- 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 an alert definition, and create a dashboard and one or more views. 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.

Verify that you are monitoring one or more vCenter Server instances. See the *vRealize Operations Manager Configuration Guide*.

Procedure

1 Troubleshoot Problems with a Host System

Use the Troubleshooting tabs to identify the root cause of problems that the system does not resolve by alert recommendations or simple analysis.

2 Examine the Environment Details

Examine the status of your objects in the views and heat maps 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.

3 Examine the Environment Relationships

Use the Environment tab to examine the status of the three badges as they relate to the objects in your environment hierarchy. You can then 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, use **All Metrics > Show Object Relationship**.

4 Fix the Problem

Use the troubleshooting features of vRealize Operations Manager to examine problems that put your objects in a critical state, and identify solutions. To resolve the resource and time remaining problems, use the Capacity Optimization function.

5 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. These tools apply the troubleshooting solutions that you used to research and solve the problems with your host system, and make the troubleshooting tools and solutions available for future use.

Troubleshoot Problems with a Host System

Use the Troubleshooting tabs to identify the root cause of problems that the system does not resolve by alert recommendations or simple analysis.

To troubleshoot the symptoms of the capacity problems that are occurring on the cluster and host system, and determine when those problems occurred, use the Troubleshooting tabs to investigate the memory problem.

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 Investigate the critical symptoms.

- a Point to 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 point to Cluster Compute Resource Time Remaining is critically low, the metric Capacity|Time Remaining 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 the USA-Cluster.

4 Click the **Events > Timeline** tab to review the triggered symptoms, alerts, and events that occurred on the 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 Point to 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 Point to the warning events.

You notice that a hard threshold violation occurred on the data center late in the evening. The hard threshold violation shows that the Badge|Workload metric value was under the acceptable value, and that the violation triggered.

- 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 the 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 following the graph.

- b Click each red triangle to identify the affected object and highlight it in that pane.

- 6 Click the **Capacity** tab to evaluate details of capacity and time remaining.

- 7 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 **All Metrics > Capacity Analytics Generated** and double-click **Capacity Remaining (%)**.

The Capacity Remaining (%) calculation appears on the right pane.

- c In the metrics pane, expand **All Metrics > Badge** and double-click **Workload (%)**. The Workload (%) calculation appears on the right pane.

- d 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 Badge|Workload (%) calculation displays workload extremes.

Results

You have analyzed the symptoms, timeline, events, and metrics related to the problems on your cluster. Through your analysis, you have determined that the heavy workload on the cluster has caused the cluster to start running out of capacity.

What to do next

Examine the Details views and heat maps to interpret the properties, metrics, and alerts. Also, look for trends and spikes that occur in the resources for your objects, the distributions of resources across your objects, and data maps. You can examine the use of various object types across your objects.

Examine the Details views and heat maps to interpret the properties, metrics, and alerts. Also, to look for trends and spikes that occur in the resources for your objects, the distributions of resources across your objects, and data maps. You can examine the use of various object types across your objects. See [Examine the Environment Details](#).

Examine the Environment Details

Examine the status of your objects in the views and heat maps 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.

Use the heat maps 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 heat map configuration.

Prerequisites

Use the Troubleshooting tabs to look for root causes. See [Troubleshoot Problems with a Host System](#).

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 the 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 Overview**, and examine the number of virtual machines listed for the 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 **Capacity** tab.
 - b In the inventory tree, expand **USA-Cluster**, and click each of the host systems in turn.
 - c The host system w2-vcopsqe2-009 is in a critical state, with no capacity remaining.
 - d Click the **Details** tab, then click **Views**, and click **Cluster Configuration View**.
 - e To reclaim capacity from several virtual machines, select the cluster name
 - f Click the **Action** menu next to the cluster, and select **Set CPU Count and Memory for VM**.

- g In the workspace that appears, 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 suggests fewer CPUs for each virtual machine.

- h Click the check box next to each virtual machine that has a suggested lower CPU count, and click **Begin Action**. A confirmation message indicates that the action is underway and provides the task ID that you use to track the action in the Recent Tasks section under Administration. 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 heat maps for the host system and virtual machine objects in the USA-Cluster.

- a In the inventory tree, click the **USA-Cluster**.
- b Click **Details**, click **Heatmaps**, and click through the list of heat map views.
- c Click **Which VMs currently have the highest CPU demand and contention?**

The heat map displays blocks that represent the objects in the USA-Cluster. The block for a virtual machine appears in red, which indicates that it has a critical problem.

- d Point to 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 heat maps to evaluate the status of your objects and identify trends and spikes, and free up capacity for your host system and the USA-Cluster. To further narrow in on problems, you can examine the other views and heat maps. You can also create your own views and heat maps.

What to do next

Examine the status for the objects in your environment hierarchy to determine which objects are in a critical state. Then examine the object relationships to determine whether a problem on one object is affecting one or more other objects.

Examine the status for the objects in your environment hierarchy to determine which objects are in a critical state. Then examine the object relationships to determine whether a problem on one object is affecting one or more other objects. See [Examine the Environment Relationships](#).

Examine the Environment Relationships

Use the Environment tab to examine the status of the three badges as they relate to the objects in your environment hierarchy. You can then 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, use **All Metrics > Show Object Relationship**.

As you click each of the badges in the Environment tab, you see that several objects are experiencing critical problems with health. Others are reporting critical risk status.

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 heat maps. See [Examine the Environment Details](#).

Examine the status of your objects in views and heat maps. 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 the **Environment** tab.
 - b On the Badge toolbar, click through the three badges - Health, Risk, and Efficiency - and look for red icons to identify critical problems.

As you click through the badges, you notice that your vCenter Server and other top-level objects appear to be healthy. However, you see that a host system and several virtual machines are in a critical state for health, risk, and efficiency.
 - c Point to 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 the **Environment** tab.
 - b Examine the badge states for the objects in USA-Cluster.
 - c Many of the objects display critical states for 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 to act on.
- 4 Examine the relationship map.
- a Click **All Metrics > Show Object Relationship**.
 - 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 following. Use the descendant object list identify all the objects related to the host system that might be experiencing problems.

What to do next

Use the user interface to resolve the problems.

Use the user interface to resolve the problems. See [Fix the Problem](#) .

Fix the Problem

Use the troubleshooting features of vRealize Operations Manager to examine problems that put your objects in a critical state, and identify solutions. To resolve the resource and time remaining problems, use the Capacity Optimization function.

You have used the Alerts, Details, All Metrics, and Environment areas of the user interface to examine critical problems such as resource contention and time remaining issues that occur on your objects. To resolve those problems, you can use the Capacity Optimization function.

Prerequisites

Examine the environment relationships. See [Examine the Environment Relationships](#).

Examine the environment relationships. See [Examine the Environment Relationships](#).

Procedure

- 1 In the menu, click **Home**, then click **Overview** under Optimize Capacity in the left pane. The Capacity Overview screen appears.
- 2 **Select** the data center - DC-Denver-19 - that contains the problem objects.

The data in the lower half of the screen refreshes to display time remaining information and reclaim recommendations for selected data center DC-Chicago-12. NOTE: Double-clicking the data center graphic displays the Object Details page for that data center.

- 3 At the graph, select **Most Constrained** from the **Sort By:** choices and **CPU** from CPU|Memory|Disk Space above the graph.

The graph refreshes to show the usage value almost touching 100% and the timeline/projection value nearly intersecting the usage value. The data center is almost out of CPU.

- 4 Scroll down the page to the Recommendations below the graph.

Option 1 lists total resources (CPU, memory, disk space) that can be reclaimed. Option 2 lists the hardware to purchase to increase time remaining to 150 days.

- 5 Click **RECLAIM RESOURCES**.

The Reclaim screen appears, displaying data for DC-Chicago-12. The "How much can you save?" pane shows that \$4140/month can potentially be saved. Looking to the top of the table, you see that the \$4140 sum appears next to Oversized VMs.

- 6 Click **Oversized VMs**. Then click the chevron next to a cluster name on the left of the table.

All the VMs in the cluster are listed.

- 7 Select the check box next to VM Name in the table heading.

All the VMs in the cluster are checked.

- 8 Click **RESIZE VM(s)**.

The Resize VMs page appears, showing the 20 VMs available for resizing.

- 9 Leave the recommendation as is, without editing the target reductions, then select the "I understand that workloads may be interrupted..." check box and click **RESIZE VM(s)**.

The system runs the resize action.

Results

You have used Capacity Optimization to resolve problems on a host system that is experiencing critical problems. The data center does not run out of CPU, and instead realizes projected cost savings of nearly \$50,000 annually.

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, configure the Workload Optimization alerts to be automated. See the *vRealize Operations Manager Configuration Guide*.

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. These tools apply the troubleshooting solutions that you used to research and solve the problems with your host system, and make the troubleshooting tools and solutions available for future use.

To 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 Cluster Utilization dashboard to view the use index, CPU demand, and memory use for your clusters. This dashboard also tracks Internet use and disk I/O operations.
- Use the Capacity Summary dashboard to track total environment capacity, system-wide capacity and time remaining, and capacity remaining by CPU, memory, and storage. The dashboard also includes Top 10 lists for clusters running out of CPU, memory, and storage, respectively. Additional details are available.
- Use the Capacity Optimization dashboard to examine the provisioned capacity levels for CPU, disk, and memory and to review potential reclaimable capacity from CPUs, data centers, snapshot waste, and virtual memory.

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 . By running the many instances, you can dedicate a monitor to each dashboard and visually track the status of your objects.

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 **Self Troubleshooting** dashboard, and review the widgets included on it: Object Type, Select Objects, Metric Picker, and Metric Chart.

By adding the Object List, Alert List, Heatmap, and Top-N widgets, you can easily peruse the status of the host systems that you select in the Object List widget. Configure widget interaction 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 **System Health**, 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 Add the Heatmap and Top_N widgets.

- f 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**.
- g 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 monitor continuously all the objects in your environment. To plan for growth and new projects, see [Chapter 7 Capacity Optimization for Your Managed Environment](#) To generate reports, and create and customize policies, see the *vRealize Operations Manager Configuration Guide* .

Troubleshooting Workbench Home Page

The **Troubleshooting Workbench** home page is where you find active troubleshooting sessions and recent searches. The active troubleshooting sessions do not persist after you log out from vRealize Operations Manager .

Where You Find the Troubleshooting Workbench Home Page

- Navigate to the **Troubleshooting Workbench** home page from **Home > Troubleshoot > Workbench**.
- From the Quick Start page, click **Workbench** in the **Troubleshoot** section.

The **Troubleshooting Workbench** home page displays a search bar, a list of active troubleshooting sessions, and recent searches. You can open a session to find potential evidences for your problems.

How Troubleshooting Workbench Home Page Works

All troubleshooting workbench sessions that are active in the current login are displayed in the **Active Troubleshooting** section of the **Troubleshooting Workbench** home page. Changes that you make to the scope, time, or potential evidences in the troubleshooting workbench page are not be saved on logging out. The next time you log in to vRealize Operations Manager , the sessions that were earlier under **Active Troubleshooting** are displayed under **Recent Searches**.

Discovering Potential Evidences Using the Troubleshooting Workbench

The Troubleshooting Workbench is where you perform advanced troubleshooting tasks on an alert that triggered on an object. You can investigate both known and unknown issues in vRealize Operations Manager .

Where You Find the Troubleshooting Workbench

You can start the Troubleshooting Workbench with an alert in context from the alert information page, or you can search for an object and start the Troubleshooting Workbench to investigate known or unknown issues related to the object.

- To start the Troubleshooting Workbench with an alert in context, in the menu, click **Alerts**. Click an alert from the alert list and click **Launch Workbench** from the **Potential Evidence** tab.
- To start the Troubleshooting Workbench with an alert in context, in the menu, click **Environment**, then select a group, custom data center, application, or inventory object. Click the object and then the **Alerts** tab. Click **Launch Workbench** from the **Potential Evidence** tab.
- To investigate known or unknown issues with an object in context, search for the object or click **Environment** to locate the object and click **Troubleshoot** on the top.

How the Troubleshooting Workbench Works

You look for potential evidences of a problem within a specific scope and time range. The **Selected Scope** control on the left of the Troubleshooting Workbench page is where you vary the scope. You can vary the scope in the following ways:

- You can select only the object that you are investigating, or include several upstream and downstream relationships by increasing the scope. As you increase the scope, more objects are displayed in the inventory tree.
- You can select a custom scope to include objects of your choice. Click **Custom** to open an interactive window where you use the pointer to visually rearrange your objects, view relationships and add peers to modify the relationships. To see details about the object, place the pointer for a few seconds above the object. You can reset a custom scope to start all over again.
- You can use the drop-down menu to narrow down the type of objects displayed.

The default time range is two hours, and thirty minutes before the alert triggered when the context is alert based, or one hour before the current time, when the context is object based. You can select a different time range, up to seven days, using the date and time controls.

The potential evidences are based on Events, Property Changes, and Anomalous Metrics which are displayed on the right of the Troubleshooting Workbench change in the **Potential Evidence** tab. Information in these sections is displayed as cards.

Events

Displays events, based on a change in the metrics. Events for metrics that have breached the usual behavior, and major events that have occurred within the selected scope and time are displayed. The cards are based on dynamic thresholds for a metric, which is calculated based on historical and incoming data.

Property Changes

Displays important configuration changes that occurred within the selected scope and time. Both single and multiple property changes are displayed. For multiple property changes, you can view the latest and previous changes.

Anomalous Metrics

Metrics which have shown drastic changes within the selected scope and time. Ranks the results based on the degree of change. The most recent anomalous metric based on a time-sliced comparison in the current time range is given the highest weightage.

You can explore more details about any of the cards displayed in the Troubleshooting Workbench by clicking the card pop-out option. You can close a card and it is no longer displayed in the Troubleshooting Workbench. To load the cards again, click **Go** in the **Time Range**.

When you pin a metric, it appears in the **Metrics** tab of the Troubleshooting Workbench. You can perform further investigation on the metric in the Metrics tab. You can compare the pinned metrics with other metrics displayed in the tab. You can close the pinned metrics and browse other metrics for specific objects.

Similarly, the **Alerts** and **Events** tabs are where you investigate the potential evidences further. You can filter and group alerts. If you want to focus on the alerts for a specific object in your selected scope, you can clear all the alerts and then click the object in the scope.

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 can take ownership of an alert or assign alerts to other vRealize Operations Manager users.

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 shortly, 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 problem 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 message event, or metric event are canceled. You cannot manually cancel other types of symptoms. If a message event symptom or metric event symptom triggered the event, then the alert is effectively canceled. If a metric symptom or property symptom triggered the alert, 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 canceled 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 you are responsible for.

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 the symptoms are triggered.

As the alerts are generated, you must process the alerts based on the negative effect they have on objects in your environment. To do the processing, 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 a memory contention that a few virtual machines cause.
- Cluster has many virtual machines that have a 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 then 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 currently 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.
- 6 To assign the ownership of the `Virtual machine has unexpected high CPU workload` alert to another user, click the alert, click **Actions** on the menu bar, and click **Assign to**.

- 7 Enter the name of user to whom you want to assign the ownership of the alert and click **Save**.

The Assigned to: field in Alert Details updates with the name of the user you have assigned the alert to.

Note You can remove the ownership assigned to a user by clicking the alert and selecting the **Release Ownership** option from the **Actions** menu.

- 8 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. Then click **Actions** on the menu bar and click **Suspend**.

- a To suspend the alert for an hour, enter **60**.

- 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.

- 9 Select the row that contains the `Cluster has many Virtual Machines that have memory contention due to memory compression, ballooning or swapping` alert. Then 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 this action does not resolve the underlying condition. 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 an Alert in the Health Alert List](#).

User Scenario: Respond to an Alert in the Health Alert List

In this scenario, you investigate and resolve the `Virtual machine has an unexpected high CPU workload` alert. The alert might be generated for more than one virtual machine.

Prerequisites

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.

- Process and take ownership of the alerts you 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 the section *Working with Actions That Use Power Off Allowed* in the vRealize Operations Manager Information Center.
- Process and take ownership of the alerts you 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* section in *vRealize Operations Manager Configuration Guide*.

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 an unexpected high CPU workload** alert name.
- 5 Review the information. To show the recommendations, click **Configuration > Recommendations** in the left pane.

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 can resolve the alert?
What is Causing the Issue?	<p>Do the triggered symptoms support the recommendations? Do the other triggered symptoms contradict the recommendation, indicating that you must investigate further?</p> <p>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.</p>
Non-Triggered Symptoms	<p>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.</p> <p>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?</p>

- 6 To resolve the alert based on the recommendation to check the guest applications to determine whether a 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 Start 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 suggested 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 To allow power off or to create a snapshot, depending on how your virtual machines are configured, select the following options.

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 any regard for the state of the operating system. In addition to the question whether the action shuts down or powers off a virtual machine, you must also 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 suggested number of CPUs to the target virtual machine.

- 8 Allow several collection cycles to run after implementing the suggested 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.

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**, **Capacity**, 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. 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 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 Health, 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 effect 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 must be addressed shortly.
- Efficiency alerts tell you where you can improve performance or reclaim resources.

Alerts for an Object or an Object Group

For a single object, the Top alerts are the alerts generated for the object. 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.

Object groups can include one object type, such as hosts, or multiple objects types, such as hosts, virtual machines, and datastores. When you are working with object groups, 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. If the average health is above the Warning, Immediate, or Critical threshold, a group population alert considers the health of all group members and is triggered. If a group population alert is generated, the alert affects the badge score and color. 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. The alerts 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 the alerts on the **Summary** tab, and resolve the problems that the alert indicates.
- 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. Using this method, 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. 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.

- Health alert:Host has memory contention caused by a few virtual machines.
- Risk alert:Virtual Machine has a chronic high memory workload.
- Risk alert:Virtual Machine is demanding more CPU than the configured limit.
- Efficiency alert:Virtual Machine has large disk snapshots.

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 *vRealize Operations Manager Configuration Guide*.
- 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 the *vRealize Operations Manager Configuration Guide*.
- 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.

- 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 a 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 a 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 in the following table.

- 9 Click each virtual machine alert and evaluate the information provided on the **Alerts** tab.

Virtual Machine Alert	Evaluation
Virtual Machine has a 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 can allow you to increase the CPU count and resolve the virtual machine alert, and might resolve the host memory contention alert.

- 10 Take the suggested actions.

Results

Your actions might resolve 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.

- You can also click **Administration > Inventory > Select an Object from the List > click Show Detail**.
- In the menu, select **Alerts** to display the All Alerts screen. Click an **alert** to display 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

The screenshot displays the VMware vSphere Summary Tab for the object `vc_10.27.83.18`. The interface includes a top navigation bar with tabs for Summary, Alerts, Metrics, Capacity, Compliance, Events, and more. The Summary tab is selected, showing a summary of the object's properties: Cluster: 1, ESXi: 4, Virtual Machine: 32, and Datastore: 5. Below this, there are sections for Consumer (Virtual Machines: 24 Running of 32) and Provider (Usable Capacity: ESXi Hosts: 4 Running of 4). At the bottom, there are two tables: one for Cluster Name, Host, Virtual Machine, Capacity Remaining, Time Remaining, and VM Remaining; and another for Datastore Name, Capacity, Virtual machine, Capacity Remaining, and Time Remaining.

Cluster Name	Host	Virtual Machine	Capacity Remaining	Time Remaining	VM Remaining
ESO-EVN-Cluster1	4	32	20	52.29 Week(s)	

Datastore Name	Capacity	Virtual machine	Capacity Remaining	Time Remaining
datastore69	1.81 TB	7	85.98 %	52.29 Week(s)
datastore42	923 GB	10	38.72 %	52.29 Week(s)
Datastore_iSCSI	14.5 GB	0	93.76 %	52.29 Week(s)
datastore37	923 GB	11	53.31 %	52.29 Week(s)
datastore59	924 GB	4	37.82 %	52.29 Week(s)

Table 6-1. Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. It also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Consumer	Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory.
Provider	Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage.

Table 6-1. Summary Tab Options (continued)

Option	Description
Cluster	Displays the cluster details of the selected object.
Datastore	Displays the datastore details of the selected object.

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 6-2. Datastore Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>

Table 6-2. Datastore Summary Tab Options (continued)

Option	Description
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available.
Performance	This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart.
Configuration	This widget displays the configuration details for the selected datastore object.

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

The screenshot displays the Host Summary tab for a VMware host. The top navigation bar includes tabs for Summary, Alerts, Metrics, Capacity, Compliance, Logs, Events, and more. The main content area is divided into several sections:

- Host Overview:** Shows the host name (evn1-hs1-0802.eng.vmware.com), status (Powered On), and hardware details (Dell Inc. PowerEdge R630, Version: 6.5.0, 10719125, CPU: 44 Cores, 96.8 GHz, Memory: 383.91 GB).
- Active Alerts:** A table showing alert levels (Critical, Immediate, Warning, Info) and their status (Self, All).
- Time Remaining:** A widget showing the host is powered on for > 1 Year.
- Capacity Remaining:** A widget showing 10% (1.83 TB) of capacity remaining, most constrained by Disk Space Demand.
- Utilization:** A table showing various utilization metrics:

Metric	Value
CPU Usage	11.15 %
Memory Usage	83.74 %
Memory Balloon	0 KB
Disk Total IOPS	1,078.87
Disk Total Throughput	10.5 MBps
Network Usage Rate	13.75 MBps
- Performance:** A table showing performance metrics:

Metric	Value
Worst Consumer CPU Ready	0.046
Worst Consumer Memory Contention	0
Worst Consumer Disk Latency	13.41
Packets Dropped	0 %
Consumers with Memory Contention	0
Consumers with CPU Ready	0
- Configuration:** A table showing configuration details:

Category	Value
Hardware	Service Tag: HZW4NK2
CPU	CPU Model: Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz
Network: Logical	Management Address: 10.27.81.2
Storage: Path	Total number of Active Path: 8

Table 6-3. Host Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available.
Performance	This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart.
Configuration	This widget displays the hardware, CPU, and Network configuration details of 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

The screenshot displays the VMware vSphere VM Summary tab for a virtual machine named **vRLI_ESO_1_small**. The interface includes a top navigation bar with tabs for Summary, Alerts, Metrics, Capacity, Compliance, Logs, Events, and more. A **TROUBLESHOOT** button is visible in the top right corner.

VM Details:

- Powered On:** Indicated by a green checkmark.
- OS:** SUSE Linux Enterprise 11 (64-bit)
- IP Address:** 10.27.74.145, 10.27.74.148
- VMware tools:** Tools Version 10.2.0, Running
- Disk Space:** 530.5 GB
- Number of virtual CPUs:** 4
- Memory:** 8 GB

Active Alerts:

Alert Type	Self	All
Critical	1	1
Immediate	0	0
Warning	0	0
Info	0	0

Time Remaining: 0 Days. Most constrained by Memory Demand.

Capacity Remaining: 0 % (0 KB). Most constrained by Memory Demand.

Utilization:

Metric	Value
CPU Usage	4.67 GHz
Free Memory	267.42 MB
Guest Page In Rate per second	74.8
Virtual Disk Total IOPS	33.93
Virtual Disk Total Throughput	543.67 KBps

Performance:

Metric	Value
CPU Ready	0.076 %
CPU Co-stop	0 %
Memory Contention	0 %
Virtual Disk Total Latency	4.88 ms
Network Transmitted Packets Dropped	0

Configuration:

- Virtual Hardware:** CPU: 4 (4 Sockets x 1 vCore)
- Resource Allocation:** CPU: No Limit, No Reservation
- Tools:** Version: 10.2.0, Guest Tools Unmanaged, Guest Tools Running
- Network:** IP Addresses: 10.27.74.145, 10.27.74.148, 00:50:56:a6:11:19
- Guest OS Partition:** /storage/core: 482.31 GB Configured, 467.55 GB Used
- Virtual Disk:** Hard disk 1: 20 GB

Table 6-4. VM Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available.

Table 6-4. VM Summary Tab Options (continued)

Option	Description
Performance	This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart.
Configuration	This widget displays the virtual hardware, resource allocation, tools, and Network configuration details of the virtual machine.

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

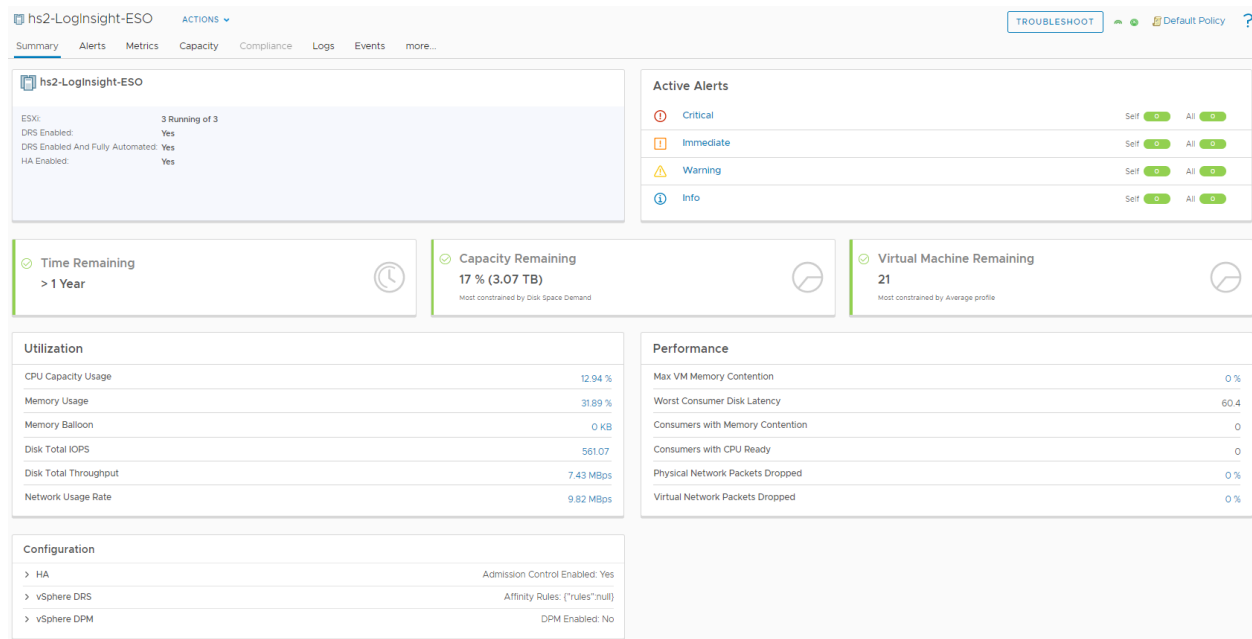


Table 6-5. Cluster Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.

Table 6-5. Cluster Summary Tab Options (continued)

Option	Description
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Virtual Machine Remaining	This widget displays the remaining virtual machines in the cluster. To see the details of the remaining virtual machines, click the Virtual Machine Remaining card.
Utilization	This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available.
Performance	This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart.
Configuration	This widget displays the configuration details of the cluster.

Data Center Summary Tab

The data center Summary tab provides an overview of the state of the selected data center. For the selected object, the data center Summary tab displays the alerts as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the data center and use the information to begin troubleshooting problems.

Understanding the Data Center Summary Tab

10.161.73.31 ACTIONS TROUBLESHOOT vSphere Solution's Default Policy (Mar 30, 2020 8:52:45 PM)

Summary Alerts Metrics Capacity Compliance Events Details Environment Reports less...

10.161.73.31

Cluster: 1
ESXi: 2
Virtual Machine: 33
Datastore: 5

Active Alerts

Critical	Self: 0	All: 0
Immediate	Self: 0	All: 0
Warning	Self: 0	All: 0
Info	Self: 0	All: 0

Consumer

Virtual Machines
33 Running of 33

vCPU: 63
RAM: 119.5 GB
Provisioned: 1.22 TB

Provider (Usable Capacity)

ESXi Hosts
2 Running of 2

CPU: 153.42 GHz
RAM: 125.53 GB
Storage: 1.48 TB

vSphere Distributed Switch Name	Version	Total Number of Hosts	Maximum number of Ports	Used Number of Ports
DSwitch	7.0.0	0	8	0

1 - 1 of 1 items

Cluster Name	Host	Virtual Machine	Capacity Remaining	Time Remaining	VM Remaining
FT_TEST_CLUSTER	2	33	2.67 %	3 Day(s)	1

1 - 1 of 1 items

Datastore Name	Capacity	Virtual machine	Capacity Remaining	Time Remaining
Datastore.0	499.75 GB	11	47.68 %	52.29 Week(s)
Datastore.1	499.75 GB	7	45.95 %	52.29 Week(s)
Datastore.2	499.75 GB	15	47.16 %	52.29 Week(s)

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Consumer	Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory.
Provider	Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage.
vSphere Distributed Switch Name	Displays the details of the vSphere distributed switch.
Metadata	Displays the metadata details of the data center.
Cluster	Displays the cluster details of the selected object.
Datastore	Displays the datastore details of the selected object.

Resource Pool Summary Tab

The Resource Pool Summary tab provides an overview of the state of the resources in the resource pool. For the selected resource, the Resource Pool 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 resource pool and use the information to begin troubleshooting problems.

Understanding the Resource Pool Summary Tab

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Utilization	This widget is used to find out the trends in capacity used by the selected resource pool as against the total capacity available.
Performance	This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart.
Resource Pool	This widget lists the resource pool name, cpu status, and memory status of the resources that are part of the corresponding resource pool.

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 6-6. Customer Group and Container Summary Tab Options

Option	Description
Recommended Actions	<p>This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance.</p> <p>The badges provide a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the badge.</p>

Workload Management Enabled Cluster Summary Tab

The Workload Management enabled cluster is a cluster with Kubernetes enabled, running on vSphere (also called Supervisor cluster). It hosts a type of resource pool called Namespaces. The Workload Management Enabled Cluster Summary tab provides an overview of the state of the selected cluster.

Understanding the Cluster Summary Tab

compute-cluster ACTIONS

Summary Alerts Metrics Capacity Compliance Events more...

compute-cluster
Workload Management Enabled

ESXi: 3 Running of 3
DRS Enabled: Yes
DRS Enabled And Fully Automated: Yes
HA Enabled: Yes

Active Alerts

Alert Type	Self	All
Critical	0	0
Immediate	0	0
Warning	0	0
Info	0	0

Time Remaining
81 Days
Most constrained by Disk Space Demand

Capacity Remaining
22 % (21.51 GB)
Most constrained by Memory Demand

Virtual Machine Remaining
14
Most constrained by Average profile

Utilization

Metric	Value
CPU Capacity Usage	22.06 %
Memory Usage	69.99 %
Memory Balloon	0 KB
Disk Total IOPS	1,008.13
Disk Total Throughput	6.12 MBps
Network Usage Rate	8.95 MBps

Performance

Metric	Value
Max VM Memory Contention	0 %
Worst Consumer Disk Latency	74.27
Consumers with Memory Contention	0
Consumers with CPU Ready	100
Physical Network Packets Dropped	0 %
Virtual Network Packets Dropped	0.00023 %

Configuration

Option	Value
HA	Admission Control Enabled: No
vSphere DRS	Affinity Rules: ("rules":null)
vSphere DPM	DPM Enabled: No

Namespaces

Metric	Value
Config Status	RUNNING
Current Version	v1.15.4-vsc0.0.1-34247796
Kubernetes Status	READY

Table 6-7. Workload Management Enabled Cluster Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object and whether the Workload Management is enabled or disabled.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the badge .</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.

Table 6-7. Workload Management Enabled Cluster Summary Tab Options (continued)

Option	Description
Virtual Machine Remaining	The virtual machine remaining number is based on the average profile. The virtual machine remaining numbers are calculated when you enable one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.
Utilization	<p>This widget is used to find out the trends in capacity used by a selected cluster as against the total capacity available. The key utilization indicators are:</p> <ul style="list-style-type: none"> ■ CPU Capacity Usage ■ Memory Usage ■ Memory Balloon ■ Disk Total IOPS ■ Disk Total Throughput ■ Network Usage Rate
Performance	<p>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. Click each metric to see the expanded chart.</p> <p>The key performance indicators are:</p> <ul style="list-style-type: none"> ■ Max VM Memory Contention ■ Worst Consumer Disk Latency ■ Consumers with Memory Contention ■ Consumers with CPU Ready ■ Physical Network Packets Dropped ■ Virtual Network Packets Dropped
Configuration	This widget displays the hardware, CPU, and Network configuration details of the host.
Namespaces	Lists the configuration status, current version and Kubernetes status of the namespaces in the cluster.

Namespace Summary Tab

A namespace sets the resource boundaries where vSphere Pods and Tanzu Kubernetes clusters created by using the Tanzu Kubernetes Grid Service can run. The Namespace summary tab provides an overview of the state of the selected Namespace.

Understanding the Namespace Summary Tab




Table 6-8. Namespace Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status based on the alert type.</p> <p>To see the alerts for the object, click the badge .</p>
Utilization	<p>This widget is used to find out the trends in capacity used by a selected namespace as against the total capacity available.</p> <p>The key utilization indicators are:</p> <ul style="list-style-type: none"> ■ CPU Usage ■ Consumed Memory

Table 6-8. Namespace Summary Tab Options (continued)

Option	Description
Performance	<p>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. Click each metric to see the expanded chart.</p> <p>The key performance indicators are:</p> <ul style="list-style-type: none"> ■ Worst Consumer CPU Ready ■ Worst Consumer Memory Contention ■ Consumers with Memory Contention ■ Consumers with CPU Ready
Configuration	<p>This widget displays the following configuration details about the Namespaces:</p> <ul style="list-style-type: none"> ■ Configuration status ■ Virtual Machines ■ Number of Tanzu Kubernetes clusters ■ Pods

vSphere Pod Summary Tab

vSphere Pods run containers without needing to customize a Kubernetes cluster. You can deploy vSphere Pods directly on ESXi hosts. It hosts a type of resource pool called Namespace. The vSphere Pod Summary tab provides an overview of the state of the vSphere Pods.

Understanding the vSphere Pod Summary Tab

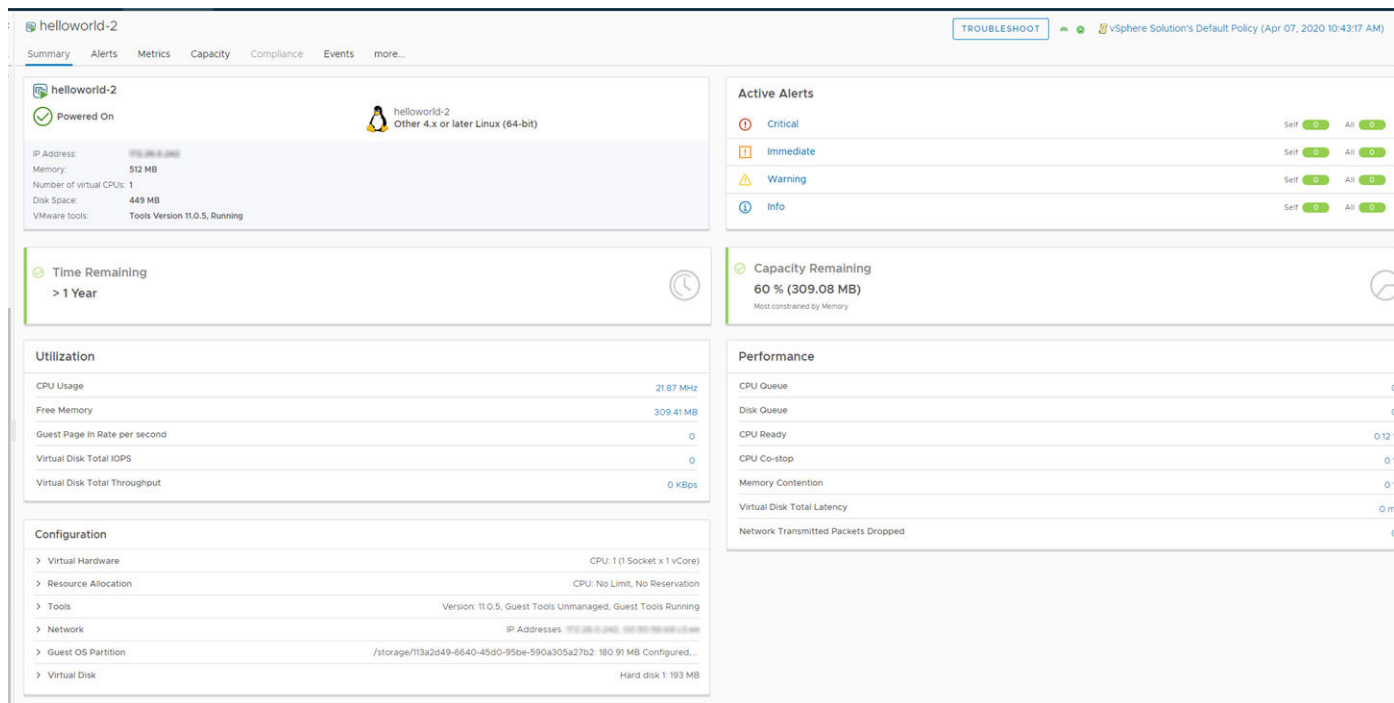


Table 6-9. vSphere Pod Tab Summary Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining until the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	<p>This widget is used to find out the trends in capacity used by a selected vSphere Pod as against the total capacity available.</p> <p>The key utilization indicators are:</p> <ul style="list-style-type: none"> ■ CPU Usage ■ Free Memory ■ Guest Page in Rate per second ■ Virtual Disk Total IOPS ■ Virtual Disk Total Throughput
Performance	<p>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. Click each metric to see the expanded chart.</p> <p>The key performance indicators are:</p> <ul style="list-style-type: none"> ■ CPU Queue ■ Disk Queue ■ CPU Ready ■ CPU Co-stop ■ Memory Contention ■ Virtual Disk Total Latency ■ Network Transmitted Packets Dropped
Configuration	This widget displays the hardware, CPU, and Network configuration details of the host.

Tanzu Kubernetes cluster Summary Tab

The Tanzu Kubernetes cluster runs Kubernetes workloads natively on the hypervisor layer. The Tanzu Kubernetes cluster Summary tab provides an overview of the state of the Tanzu Kubernetes clusters.

Understanding the Tanzu Kubernetes cluster Summary Tab




Table 6-10. Tanzu Kubernetes cluster Tab Summary Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> Health alerts that usually require immediate attention. Risk alerts indicating that you must look into any problems shortly. Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the badge .</p>

Table 6-10. Tanzu Kubernetes cluster Tab Summary Options (continued)

Option	Description
Utilization	<p>This widget is used to find out the trends in capacity used by a selected Tanzu Kubernetes cluster as against the total capacity available.</p> <p>The key utilization indicators are:</p> <ul style="list-style-type: none"> ■ CPU Usage ■ Consumed Memory
Performance	<p>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. Click each metric to see the expanded chart.</p> <p>Key performance indicators are:</p> <ul style="list-style-type: none"> ■ Worst Consumer CPU Ready ■ Worst Consumer Memory Contention ■ Consumers with Memory Contention ■ Consumers with CPU Ready

vSAN Cluster Summary Tab

The vSAN Cluster tab provides an overview of the state of the selected vSAN cluster. For the selected object, the vSAN cluster tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN cluster and use that information to begin troubleshooting problems.

Where To View vSAN Cluster Summary Page

On the menu, click **Environment > VMware vSAN > vSAN Core Services and Hardware > vSAN Cluster**.

Table 6-11. vSAN Cluster Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>

Table 6-11. vSAN Cluster Summary Tab Options (continued)

Option	Description
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected vSAN cluster as against the total capacity available.
Configuration	This widget displays the configuration details of the cluster.
Contention	This widget displays the memory contention details of the vSAN cluster.

vSAN Cluster Disk Group Summary Tab

The vSAN Cluster Disk Group Summary tab provides an overview of the state of the selected vSAN Disk Group. For the selected object, the vSAN Disk Group tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN Disk Group and use that information to begin troubleshooting problems.

Where To View vSAN Cluster Disk Group Summary

On the menu, click **Environment > VMware vSAN > vSAN and Storage Devices > vSAN Cluster > Host System > Disk Group**.

Table 6-12. vSAN Cluster Disk Group Summary Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.

Table 6-12. vSAN Cluster Disk Group Summary Options (continued)

Option	Description
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected vSAN cluster disk group as against the total capacity available.
Contention	This widget displays the memory contention details of the vSAN cluster.
Resync	This widget displays the throughput and latency details for the vSAN cluster disk group.

vSAN Capacity Disk Summary Tab

The vSAN Capacity Disk tab provides an overview of the state of the selected vSAN capacity disk. For the selected object, the vSAN capacity disk tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN capacity disk and use that information to begin troubleshooting problems.

Table 6-13. vSAN Capacity Disk Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected capacity disk as against the total capacity available.
Contention	This widget displays the memory contention details for the selected capacity disk.

vSAN Cache Disk Summary Tab

The vSAN Cache Disk tab provides an overview of the state of the selected vSAN cache disk. For the selected object, the vSAN cache disk tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN cache disk and use that information to begin troubleshooting problems.

Table 6-14. vSAN Cache Disk Summary Tab Options

Option	Description
Troubleshoot	Start the Troubleshooting Workbench with the current object in context.
Object Summary	This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object.
Active Alerts	<p>This widget provides a visual indicator of the alert status for the following alert types.</p> <ul style="list-style-type: none"> ■ Health alerts that usually require immediate attention. ■ Risk alerts indicating that you must look into any problems shortly. ■ Efficiency alerts indicating that you can reclaim resources. <p>To see the alerts for the object, click the labels of the alert.</p>
Time Remaining	This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity.
Capacity Remaining	This widget displays the unused capacity of your virtual environment to accommodate new virtual machines.
Utilization	This widget is used to find out the trends in capacity used by a selected vSAN cache disk as against the total capacity available.
Contention	This widget displays the memory contention details for the selected cache disk.

vSAN Cluster Fault Domain Summary Tab

The vSAN cluster fault domain summary tab provides details about CPU, CPU Cores, Memory, Disc Space and Alerts associated with the fault domain of the vSAN cluster.

Where To View vSAN Cluster Fault Domain Summary

On the menu, click **Environment > VMware vSAN > vSAN and Storage Devices > vSAN Cluster > Fault Domain**.

You can also view relationship details and heat map details for the selected vSAN fault domain. The relationship section provides information about the relationship between the objects in your vSAN cluster. The heat map helps you to identify potential problems for the objects in your vSAN fault domain.

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 early notification about problems in your environment before a user calls you to report a problem. As well, object alerts can 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. Conversely, you can examine how problems reflected in alerts on other objects affect the current object.

Depending on the 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 is removed from a host for replacement. Also cancel alerts that are known issues that you cannot resolve currently because of resource constraints. Canceling an alert that is generated because of only message event or metric event symptoms cancels the alert permanently. If the underlying metric or property condition remains true, canceling an alert that is generated because of metric, super metric, or property symptoms can result in the alert being regenerated . It is only effective to cancel alerts generated because of 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 PowerOff section in *vRealize Operations Manager Configuration Guide*. .
- 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 the section on Working With Actions That Use Power Off Allowed in the vRealize Operations Manager Information Center.

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 **All Metrics** tab.

The **All Metrics > Badge > Workload %** generates a graph in the right pane that shows the workload is heavy.

- 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 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 suggested 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 system groups the alerts 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 with the alert. The system offers recommendations for addressing the alert and links to additional information. A **Run Action** button might appear in the details. Point to the button to learn what recommendation is performed if you click the button. To return to the list view, click the **X** at the top right of the alert details.

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 display 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 display 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 6-15. 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	<p>Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.</p> <p>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.</p>
Delete Canceled Alerts	Delete canceled (inactive) alerts by making a group selection or by individually selecting alerts. You cannot delete active alerts.
Suspend	<p>Suspend an alert for a specified number of minutes.</p> <p>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.</p> <p>The user who suspends the alert becomes the assigned owner.</p>
Take Ownership	<p>As the current user, you make yourself the owner of the alert.</p> <p>You can only take ownership of an alert, you cannot assign ownership.</p>
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...	<p>Offers two options for disabling the alert:</p> <p>Disable the alert in all policies: this disables the alert for all objects for all the policies.</p> <p>Disable Alert in Selected Policies: this disables the alert for objects having the selected policy. This method works only for objects with alerts.</p>
Open an external application	<p>Actions you can run on the selected object.</p> <p>For example, Open Virtual Machine in vSphere Client.</p>

Table 6-16. 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 6-17. 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 6-18. 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.

Table 6-18. 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.

Table 6-19. All Filters

All Filters	Descriptions
Filtering options	Limit the list of alerts to those matching the filters you select. For example, you might have chosen the Time option in the Group By menu. Now you can select 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.

Table 6-19. All Filters (continued)

All Filters	Descriptions
Control State	<p>State of user interaction with the alert. Possible values include:</p> <ul style="list-style-type: none"> ■ Open. The alert is available for action and has not been assigned to a user. ■ Assigned. The alert is assigned to the user who is logged in when that user clicks Take Ownership. ■ Suspended. The alert was suspended for a specified amount of time. The alert is temporarily excluded from affecting the health, risk, and efficiency of the object. This state is useful when a system administrator is working on a problem and does not want the alert to affect the health status of the object.
Object Type	Type of object on which the alert was generated.
Updated On	<p>Date and time when the alert was last modified.</p> <p>An alert is updated whenever one of the following changes occurs:</p> <ul style="list-style-type: none"> ■ Another symptom in the alert definition is triggered. ■ Triggering symptom that contributed to the alert is canceled.
Canceled On	<p>Date and time when the alert canceled for one of the following reasons:</p> <ul style="list-style-type: none"> ■ 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.

Table 6-20. Alert Details Tab


Section	Description
Recommendations	View recommendations for the alert. To resolve the alert, click the Run Action button if it appears.
Other Recommendations	Collapse the section to view additional recommendations. See the links in the Need More Information? section to view additional metrics, events, or other details that appear as a link.
Symptoms	View the symptoms that triggered the alert. Collapse each symptom to view additional information.

Table 6-20. Alert Details Tab (continued)

Section	Description
Alert Information	View information such as the start time, update time, and status of the alert.
Close	Click the X icon to close the alert details tab.

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 the *Definitions for Metrics, Properties, and Alerts* Guide.

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 graph and metric charts. The relationship graph 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

You can double-click any object in the graph and view the specific parent-child objects for the focus object. If you point to an object icon, you can see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object. If you double-click an object icon, the selected object becomes the focus of the map. The graph 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 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 locate the objects that you want.

All Metrics Options

The options include the graph toolbar, the metric selector options, the metric charts toolbar, and the toolbar on each chart.

Table 6-21. Relationship Map

Option	Description
Reset to initial object	Returns the map to original object if you double-clicked on an icon to examine another object.
Vertical/Horizontal	Displays a vertical or horizontal view of the graph or tree view.
Hide Text/Show Text	Hides or displays the object names.
Standard View/Fit View	<p>The Standard View option fixes the view to a specific zoom level.</p> <p>The Fit View option adjusts the graph or tree view to fit the screen.</p>
Group Items/Ungroup Items	Groups by objects types. You can view further details by double-clicking on the object. You can also choose to display the graph or tree view without grouping the object types.
Path Exploration	Displays the relative relationship path between two selected objects on the graph or tree view. To highlight the path, click the Path Exploration icon and then select the two objects from the graph or tree view.

The chart options are used to limit the metric list.

Table 6-22. Metric Chart Selector

Option	Description
Show collecting metrics	Updates the list to display only the currently collected metrics for the object.
Show previewable super metrics	Updates the list to display super metrics for the object. Note The super metrics only appear if the super metric is associated with the object, see Create a Super Metric topic in <i>vRealize Operations Manager Configuration Guide</i> .
Actions	Click the Actions icon to configure metric groups. Verify that you hold the PowerUser or administrator role. <ul style="list-style-type: none"> ■ 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.
Time Range	Filters the metrics to show only the ones that have received data in the selected time range.
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 6-23. 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.
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.

Table 6-23. Metric Chart Toolbar (continued)

Option	Description
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 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.
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 you to begin constructing a new set of charts.

Manage individual charts with the toolbar options.

Table 6-24. 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.
Correlation	<p>Runs metric correlation using the following options:</p> <p>Self-Metrics: Runs metric correlation on all metrics for the selected object, to find metrics of similar, or opposite behavioral change for the same time period. The instanced metrics are not assessed in the self-metrics correlation method.</p> <p>Peers: Runs metric correlation on the same metrics for all peer objects, to find the same metrics with behavioral changes within peer objects. Peer objects are the direct child objects of the parent for the selected objects. The child objects have the same object type.</p> <p>Note The correlation results only appears if there are at least 11 data points and the time range is within the three months period to run the metric correlation.</p> <p>Scope: Runs metric correlation on all metrics for the selected object with the selected scope, to find metrics of similar, or opposite behavioral change for the same time period. The instanced metrics are not assessed in the scope correlation method.</p> <p>After you run the correlation, the results are displayed in the Correlation window. By default, only the first 10 results for correlated metrics are displayed. To view the full list, click Show More.</p> <p>You can zoom in to view the correlated metrics and also pin them so that they appear in the preview section of the All Metrics tab.</p> <p>Note During the correlation process, some metrics are left out. For example, the badge and vRealize Operations Manager generated metrics. By default, the instanced metrics are omitted, except those in the Aggregate of all instances group.</p>
Save a Snapshot	<p>Creates a PNG file of the current chart. The image is the size that appears on your screen.</p> <p>You can retrieve the file in your browser's download folder.</p>
Save a Full Screen Snapshot	<p>Downloads the current graph image as a full-page PNG file, which you can display or save.</p> <p>You can retrieve the file in your browser's download folder.</p>
Create an Alert Definition	Allows you to create an alert for an object type or metric in a quick and easier way. For details, see <i>Create a Simple Alert Definition</i> section in <i>vRealize Operations Manager Configuration Guide</i> .

Table 6-24. Individual Metric Charts Toolbar (continued)

Option	Description
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.
Scales	<p>You can choose a scale for a stacked chart.</p> <ul style="list-style-type: none"> ■ Select Linear to view a chart in which the Y-axis scale increases in a linear manner. For example, the Y-axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on. ■ Select Logarithmic to view a chart in which the Y-axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values. <p>Note If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph.</p> <ul style="list-style-type: none"> ■ Select Combined to view overlapping graphs for the metrics. The chart uses individual scales for each graph instead of using a relative scale, and displays a combined view of the graphs. ■ Select Combined by Unit to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs.
Move Down	Moves the chart down one position.
Move Up	Moves the chart up one position.
Close	Deletes the chart.
Vertical resize	Resizes the height of a graph in the chart.
Remove icon next to each metric name in a stacked chart	Removes the graph for the metric from the chart.

Capacity Tab Overview

Use the Capacity tab to assess workload status and resource contention in the selected object. You can determine time, capacity and VM remaining until CPU, memory, or storage resources run out. With robust capacity planning and optimization, you can manage your production capacity effectively as your organization addresses changing requirements.

Capacity Tab

The **Capacity** tab provides Time Remaining and Capacity Remaining data for the selected object. Virtual Machine Remaining data is available for Clusters, Datacenters, CDC, and VC based on the average profile, or when you enable one or more custom profiles in the policy.

Where You Find the Capacity Tab

- In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. The Object details screen appears. Click the **Capacity** tab.
- In the menu, select **Alerts** to display the **All Alerts** screen. Click an **alert** to show 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 **Capacity** tab.

Understanding the Capacity Tab

For the selected object, the **Capacity** tab lists two panes with the Time Remaining and Capacity information. These panes display the value of the resources remaining till they run out.

Below the **Time Remaining** and **Capacity** panes, the time and capacity utilization metric for CPU, memory, and disk space are displayed in three panes. By default, the most constrained resource is selected. Click **CPU**, **Memory**, or **Disk Space** to change the views to these resources. These panes display the resource information based on the Demand model (default) or Allocation model (if configured).

Time Remaining Pane

When you select the **Time Remaining** pane and click one of the resource types, the utilization graph displays the historical value of the utilization metric and its forecast plotted against time, projecting how swiftly resource utilization is approaching the usable capacity.

Capacity Pane

The **Capacity Remaining** pane indicates the unused capacity of your virtual environment to accommodate new virtual machines. vRealize Operations Manager calculates the Capacity Remaining as a percentage of the remaining capacity, compared to the total. Capacity Remaining is calculated as the utilization metric forecast 3 days from now subtracted from the Usable Capacity. vRealize Operations Manager calculates the average profile and always computes the virtual machine remaining number based on the average profile. You can change the profile by clicking the + icon above the bar chart. vRealize Operations Manager calculates virtual machine remaining numbers when you enable one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.

When you select Capacity and click one of the resource types, a bar chart and a table of values based on the Demand and Allocation model (if configured) appears. The bar chart displays total usable resource, the percentage used, the percentage allocated for high availability and buffer, and the percentage remaining based on the Demand and Allocation models (if configured).

The table displays the following information for each resource type:

- **Total:** The total usable capacity for each resource type based on the Demand model or Allocation model (if configured). The difference in Total capacity and Usable capacity is set in the HA (admission control) that is set in the clusters in vSphere.

- **Usable:** The total usable capacity for each resource type based on the Demand model or Allocation model (if configured).
- **Used:** Approximate value how much utilization do you have now. Shows the forecast value of utilization metric in 3 days from now. If Capacity Remaining is greater than zero, then $\text{Used} = \text{Usable} - \text{Capacity Remaining}$.
- **Recommended Size:** The Total Capacity that must be available for a green level of Time Remaining. The slider in the policy controls the Time Remaining green zone, and the default value is 150 days.
- **Remaining:** The Capacity Remaining metric value and also the percentage. The value of Capacity Remaining metric is calculated by forecasting the utilization metric 3 days from now and subtracting it from Usable capacity.
- **Buffer:** The percentage of the capacity buffer based on the buffer value that you set in the policy. The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required.
- **High Availability:** The percentage of the high availability based on the high availability buffer.

The **Capacity** tab is a subset of the Capacity optimization capability. For additional details, refer to [Capacity Overview](#).

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 the **Alerts** 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. To 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 display 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 display the object's **Summary** tab. Click the **Alerts > Symptoms** tabs.

Table 6-25. Symptoms Data Grid

Option	Description
Criticality	<p>Criticality is the level of importance of a symptom in your environment.</p> <p>The level is based on the same level assigned when the symptom was created. The possible values include:</p> <ul style="list-style-type: none"> ■ Critical ■ Immediate ■ Warning ■ Information
Symptom	Name of the triggered symptom.
Status	<p>Current state of the symptom.</p> <p>Possible values are Active or Inactive.</p>
Created On	Date and time when the alert was generated.
Canceled On	Date and time when the symptom was canceled.
Information	<p>Information about the triggering condition for the symptom, including the trend and current value.</p> <p>The sparkline displays a range of data that includes six hours before the symptom update time and one hour after the update time.</p>

Table 6-26. 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 time. Use the timeline to identify common trends over time that are contributing to the 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. 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 the system triggers a symptom and then cancels it 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 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 the timeline in one of the three tiers. Then 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 display 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 display the object's **Summary** tab. Click the **Events > Timeline** tabs.

Table 6-27. 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 6-28. 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 6-29. 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 datastore 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.

Table 6-29. Event Filters (continued)

Option	Description
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 behavior that differs from the expected one.

Table 6-30. 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. Use the **Events** tab to compare the occurrence of events with the generated alerts. These comparisons can help 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.

You might monitor the same virtual machines with other applications that provide information to vRealize Operations Manager, with the adapters for those applications configured to provide change events. In this instance, 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 defined by a change in the 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 display 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 display 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 6-31. 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 6-32. 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 6-33. 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 datastore 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 behavior that differs from the expected one.

Table 6-34. 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: <ul style="list-style-type: none"> ■ 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 a 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 depends 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 6-35. 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. Say, for example, 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.

For Distribution views, you can click on a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment, 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 locate quickly 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 production 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 automatically 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 group 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 a 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. Also select 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 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 Maps** 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 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 locate quickly 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 6-36. 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 automatically 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 that 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 6-37. Heat Map Configuration Options

Option	Description
Configurations	<ul style="list-style-type: none"> ■ Add a 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.

Table 6-37. Heat Map Configuration Options (continued)

Option	Description	
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 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 possible 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 data center 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 To examine the resources for the cluster or data center, click **Details** in the pop-up window .

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 data center 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 data center 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.

Workload Tab

Workload metrics measure an object's demand for resources versus the actual capacity that the object can access. Use Workload values as an investigative tool when you are researching capacity constraints or evaluating the general state of objects in your environment.

Object Workload

The Workload tab present data about a single object as follows:

- The Business Week Workload - this measure reflects the system's calculation of how much capacity an object demands over a time period. The analysis compares an object's overall average workload against its capacity for a six-week period, hour by hour. Results are color-coded to show different demand levels. See the color key that follows these descriptions.
- Workload Breakdown - Data is given for the individual resources of the workload, for example, CPU and memory. The values are recalculated every five minutes.

Custom Group Workload

The Workload tab presents information for a custom group, for example vSphere World, differently from how it presents object data:







- **Current Workload Breakdown** - the system presents workload constraints in several formats: pie chart, badges, bar chart, and grid. See the color key that follows these descriptions.

Table 6-38. Custom Workload Breakdown

Format	Content
Pie Chart	Each slice of the pie represents the percentage of total workload being occupied by objects in a given state: normal, warning, critical, and so on. Point to a slice to make the percentage appear as a tool tip.
Badges	Each colored badge represents a state and includes the number of objects in a given state, for example, immediate (attention needed). You can toggle the data between the number of objects in a given state and the percentage of objects in a given state. A caption notes the total number of objects in the group.
Bar chart	A visual presentation of the percentage of all objects experiencing workload issues during that past four weeks.
Grid	All objects in the group are listed by name, object type, current level of criticality, and general issue description. You can click any object name to view the details for that object, including its Object Workload details.

Object State Color Key

Table 6-39. Object Workload States

Badge Color	Description	User Action
	Workload on the object is not excessive.	No attention required.
	Object is experiencing some high-resource workloads.	Check and take appropriate action.
	Workload on the object is approaching its capacity in at least one area.	Check 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.
	No data is available.	
	Object is offline.	

Here is a list of metrics by which the data in the Workload Tab is represented, for all interested object types.

Table 6-40. vCenter Server

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Demand Without Overhead
CPU-Usage	CPU VM CPU usage
CPU-Reserved	CPU Reserved Capacity
CPU-Overhead	CPU Overhead
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Machine Demand
Memory-Usage	Memory Host Usage
Memory-Reserved	Memory Reserved Capacity
Memory-Overhead	Memory ESX System Usage
Memory-Entitlement	Memory Usable Capacity

Table 6-41. Datacenter

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Demand Without Overhead
CPU-Usage	CPU VM CPU usage
CPU-Reserved	CPU Reserved Capacity
CPU-Overhead	CPU Overhead
CPU-Entitlement	CPU Usable Capacity
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Machine Demand
Memory-Usage	Memory Host Usage
Memory-Reserved	Memory Reserved Capacity
Memory-Overhead	Memory ESX System Usage
Memory-Entitlement	Memory Usable Capacity

Table 6-42. Cluster Compute Resource

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Demand Without Overhead
CPU-Usage	CPU VM CPU usage
CPU-Reserved	CPU Reserved Capacity
CPU-Entitlement	CPU Usable Capacity
CPU-Overhead	CPU Overhead
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Machine Demand
Memory-Usage	Memory Host Usage
Memory-Reserved	Memory Reserved Capacity
Memory-Entitlement	Memory Usable Capacity
Memory-Overhead	Memory ESX System Usage

Table 6-43. Host System

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Demand Without Overhead
CPU-Usage	CPU VM CPU usage
CPU-Reserved	CPU Reserved Capacity
CPU-Overhead	CPU Overhead
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Machine Demand
Memory-Usage	Memory Host Usage
Memory-Reserved	Memory Reserved Capacity
Memory-Overhead	Memory ESX System Usage

Table 6-44. Virtual Machine

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Demand
CPU-Usage	CPU Usage
CPU-Limit	CPU Effective limit
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Utilization
Memory-Usage	Memory Guest Usage
Memory-Reserved	Memory Reservation Used
Memory-Limit	Memory Effective limit

Table 6-45. Resource Pool

Data	Metric Name
CPU-Capacity	CPU Total Capacity
CPU-Demand	CPU Usage
CPU-Usage	CPU Usage
CPU-Reserved	CPU Reservation Used
Memory-Capacity	Memory Total Capacity
Memory-Demand	Memory Guest Demand
Memory-Usage	Memory Consumed
Memory-Reserved	Memory Reservation Used

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

When you select an object from the inventory of your environment and display the Object Details screen, you can display an overview of the related objects by clicking the Environment tab. The tab shows all the objects in your environment that are related to the selected object, with a status badge for each object. Use the Environment tab to identify related objects in your environment with health, risk, or efficiency problems.

Example: Use the Environment Tab to Find Problems

Suppose that you are trying to investigate the reason for slow performance in the 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

Now you can investigate the reason for the problem. For example, once it is determined whether the problem is chronic or temporary, you can decide how to address it. See [Using Troubleshooting Tools to Resolve Problems](#).

Environment Objects 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 Tab 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.

Where You Find the Environment Objects Tab

- In the menu, click **Environment**, then **click** a group, custom data center, application, or inventory object to display the Object Summary screen. Click the **Environment tab**.
- Alternatively, click **Environment**, then use the hierarchies in the left pane to click down to the object you want. **Click** the object to display the Object Summary screen, then click the **Environment tab**.

Table 6-46. 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	<p>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 Actions" in the <i>vRealize Operations Manager Configuration Guide</i>.</p> <p>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 Actions" in the <i>vRealize Operations Manager Configuration Guide</i>.</p>
Sort	Changes the order in which the objects are listed. Alphabetical sort is by object name.

User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options

One of your customers reports poor performance for a 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 a virtual machine, sales-10-dk. The reported conditions are poor application performance, including slow load times and slow boot, some applications are taking longer and longer to load, and files are taking longer to save. Today applications 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. The triggered symptoms indicate memory stress and the recommendation is to add more memory.

Based on experience, you are not convinced that this alert indicates the root cause, so you review the **Capacity** tab. The **Capacity** tab 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 **Capacity** tab 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 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 alerts can 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. You identified odd behavior related to a guest file system that you want to examine in the context of other metrics. This investigation can determine whether you 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, 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 the event types.
Consider whether any changes correspond to other events.
- 5 Click **View From > Parent > Select All** and click through the alerts 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 alerts 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 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.

When you are troubleshooting problems, you can run the actions from the center pane Actions menu. Alternatively, you can run them from the toolbar on list views that contain the supported objects.

When an alert is triggered, and you determine that the suggested 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 Serve Cloud Account* in *vRealize Operations Manager Configuration Guide*.
- Verify that the vCenter Adapter is configured to run actions for each vCenter Server instance. See the *vRealize Operations Manager Configuration Guide*.
- 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* section in *vRealize Operations Manager Configuration Guide*.
- 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 the section *Working With Actions That Use Power Off Allowed* in the vRealize Operations Manager Information Center.

Procedure

- 1 Select the object in the Environment page inventory trees or select one or more objects 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 To run the action on the object, select the check box 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 suggested 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 data center or data center, 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 two data centers are experiencing extreme differences in workload - one high and one low - 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 pressure 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 suggested list provided by the action to balance the workload. If you do not want to act on the entire set of objects to resolve the problem with workload, 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, because the vApp can 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 data center has unbalanced workload
- Data center 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 data center 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, view the workload on the data center to confirm that the workload was rebalanced and that the alert is gone.

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 suggested 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 list when you click **Administration** in the menu, click **Inventory**, then 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.
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 selects the host automatically.
Reason	Describes the action to be taken and the reason why the move is suggested. 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 6-47. 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 those 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 those virtual machines that are powered on, but that are in an 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 list click **Administration** in the menu, then click **Inventory**, then 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 6-48. 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 list when you click **Administration** in the menu, click **Inventory**, then 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 6-49. 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 .

Execute Script Action

To troubleshoot particular processes, you can upload a script or run a command to receive specific information. You can view the standard output or standard error as applicable.

Where You Run the Action

For supported objects and object levels, in the main menu, select the **Environment** tab and then select the relevant VM from the Inventory tree. This action is available from the **Actions** menu just below the top menu in vRealize Operations Manager .

Prerequisites

- VMware Tools must be installed and running on the VM. For details see [KB 75122](#)
- Service discovery is enabled with the successful discovery of VMs.
- The VM must be powered on and connected.

Action Options

Enter the VM credentials to authenticate even when the VM guest OS authentication status is "Success". You can run a script by entering it directly or by uploading a script file by optionally providing arguments.

Option	Description
Upload File	Use this option to browse and upload the script that you want to run.
File	Browse and upload the script file.
Args	List the arguments in the script.
Command	Select the option and enter a command in the text box.
Timeout	Script execution timeout on VMs. Script execution continues even if the dialog box is closed. You can verify the status from Administration > History > Recent Tasks .
Execute	Runs the script or command.
stdout	Displays the standard output.
stderr	Displays errors, if any.

Get Top Processes Action

The Get Top Processes action is used for troubleshooting process issues and resource issues related to the applications of the virtual machine.

How the Action Works

The Get Top Processes action, provides the status of top 10 processes for the selected virtual machine. You can troubleshoot issues related to the resources that are affecting the applications in the virtual machine.

By default, the details of top 10 processes are displayed for the selected virtual machine. You can change the number of processes and view the details for top N processes where N is between 1-100. You have the option to view the processes based on CPU and Memory.

The Get Top Processes action is run on both Windows virtual machine and Linux virtual machine. You can view the summary information for the commands only in a Linux virtual machine.

Where You Run the Action

For supported objects and object levels, in the main menu, select the **Environment** tab and then select the relevant VM from the Inventory tree. This action is available from the **Actions** menu just below the top menu in vRealize Operations Manager .

Prerequisites

- VMware Tools must be installed and running on the VM. For details see [KB 75122](#)
- Service discovery is enabled with the successful discovery of VMs.
- The VM must be powered on and connected.

Action Options

You must enter the VM credentials to authenticate when the VM is monitored in a credential-less mode or when the VM is monitored in a credential-based mode where the user is not authenticated. To ensure that you are taking the right action, review the following information.

Option	Description
Number of Processes	Displays the number of processes for which the details are displayed.
Refresh	Displays new data about processes, when you change the value for the number of processes.
Command	Displays the name of the application
PID	Displays the process ID.
CPU	Displays the CPU usage in percentage for Linux VMs. Displays the CPU usage in seconds for Windows VMs. The count starts when you start the operating system in the VM .
Mem (%)	Displays the Memory usage in KB.
User	Displays the user name.
Status	Displays the process status. It can be in one these states: <ul style="list-style-type: none"> ■ For Linux - I, R, S ■ For Windows - Unknown, Running, and Sleeping
Run	Displays data about the specified numbers of processes.

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 many 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 are violated by the Move VM action.

Important Do not attempt to move virtual machines that are members of a vApp, because the vApp can 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 requires 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. As a result, 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 list when you click **Administration** in the menu, click **Inventory**, then 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 fits 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 is 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. <ul style="list-style-type: none"> ■ Empty. vSphere DRS rules are not defined. ■ Green check mark. The move of virtual machines does not violate affinity rules. ■ Red circle with bar. The move of virtual machines does 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.

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

Table 6-50. 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 shut down 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 does 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

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 list when you click **Administration** in the menu, click **Inventory**, then 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: <ul style="list-style-type: none"> ■ 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 <code>cpu_usage_average</code> .
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 6-51. 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

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 list when you click **Administration** in the menu, click **Inventory**, then 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 so you can be sure you are taking the right action.

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: <ul style="list-style-type: none"> ■ 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 6-52. 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

To start one or more virtual machines that are in a powered off state, use the Power On VM action. 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. The action does not affect virtual machines that are currently powered on.

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

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 list when you click **Administration** in the menu, click **Inventory**, then 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

To ensure that you are taking the right action, review the following information .

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
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 6-53. 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

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 list when you click **Administration** in the menu, click **Inventory** , then 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

To ensure that you are submitting the action for the right objects, review the following information.

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.
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 6-54. 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 perform several tasks. The action determines the power state of the target virtual machines, takes a snapshot when you request it and powers off the machine if necessary and you request it. As well, the action 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 is installed, then the virtual machines are shut down before they are powered off.

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 list when you click **Administration** in the menu, click **Inventory**, then 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	<p>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.</p> <p>If you modify a value, the check box is selected. The check box must be selected to enable the OK button.</p>
Name	Name of the virtual machine as it appears in the environment inventory.
New CPU	<p>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.</p> <p>The value that appears is the calculated suggested 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 six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.</p>
Current CPU	Number of configured CPUs.

Option	Description
Power State	Indicates whether the virtual machine is powered on or powered off.
Power Off Allowed	<p>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.</p> <p>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.</p> <p>See Working with Actions That Use Power Off section in <i>vRealize Operations Manager Configuration Guide</i>. .</p>
Snapshot	<p>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.</p> <p>The name of the snapshot is supplied in the Recent Tasks messages for the action.</p> <p>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.</p>
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 6-55. Task ID Dialog Box

Option	Description
OK	To close the dialog box without further action, click OK .
Recent Tasks	To view the status of the job and verify that the job finished, click Recent Tasks .

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 you 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

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 list when you click **Administration** in the menu, click **Inventory**, then 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

To ensure that you are submitting the action for the right objects, review the following information.

Option	Description
Selected objects	<p>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.</p> <p>If you modify a value, the check box is selected. The check box must be selected to enable the OK button.</p>
Name	Name of the virtual machine as it appears in the environment inventory.
New Resv (KB)	<p>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:</p> <ul style="list-style-type: none"> ■ 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 configured as the guaranteed memory for the virtual machine.
New Limit (KB)	<p>Maximum amount of memory in kilobytes that the virtual machine can consume when the action is completed.</p> <p>The limit supports the following possible values:</p> <ul style="list-style-type: none"> ■ 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. ■ If 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 6-56. 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 CPU Count for Virtual Machine Action

The Set CPU action 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. If you are decreasing the CPU count, the action is required. 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 is 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.

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 list when you click **Administration** in the menu, click **Inventory**, then 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	<p>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.</p> <p>If you modify a value, the check box is selected. The check box must be selected to enable the OK button.</p>
Name	Name of the virtual machine as it appears in the environment inventory.
New CPU	<p>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.</p> <p>The value that appears is the calculated suggested 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 six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.</p>
Current CPU	Number of configured CPUs.
Power State	Indicates whether the virtual machine is powered on or powered off.
Power Off Allowed	<p>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.</p> <p>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.</p> <p>See Working with Actions That Use Power Off section in <i>vRealize Operations Manager Configuration Guide</i>. .</p>

Option	Description
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 6-57. Task ID Dialog Box

Option	Description
OK	To close the dialog box without further action, click OK .
Recent Tasks	To view the status of the job and verify that the job finished, click Recent Tasks .

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 configure with the current value.

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 list when you click **Administration** in the menu, click **Inventory**, then 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

To ensure that you are submitting the action for the right objects, review the following information.

Option	Description
Selected objects	<p>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.</p> <p>If you modify a value, the check box is selected. The check box must be selected to enable the OK button.</p>
Name	Name of the virtual machine as it appears in the environment inventory.
New Resv (MHz)	<p>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).</p> <p>The reservation supports the following possible values:</p> <ul style="list-style-type: none"> ■ If you set the value to 0, the virtual machine is allocated only the 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 configured as the guaranteed CPU resources for the virtual machine.
New Limit (MHz)	<p>Maximum amount of CPU consumption in megahertz that the virtual machine can consume when the action is completed.</p> <p>The limit supports the following possible values:</p> <ul style="list-style-type: none"> ■ 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.
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 6-58. 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 .

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 off 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. The action also creates a snapshot when requested and changes the number of vCPUs and memory based on the new CPU count and memory values you provided. As well, the action 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 configure with the current value.

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 list when you click **Administration** in the menu, click **Inventory**, then 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	<p>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.</p> <p>If you modify a value, the check box is selected. The check box must be selected to enable the OK button.</p>
Name	Name of the virtual machine as it appears in the environment inventory.
New CPU	<p>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.</p> <p>The value that appears is the calculated suggested 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 six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric.</p>
Current CPU	Number of configured CPUs.
Power State	Indicates whether the virtual machine is powered on or powered off.
Power Off Allowed	<p>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.</p> <p>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.</p> <p>See <i>Working with Actions That Use Power Off</i> section in <i>vRealize Operations Manager Configuration Guide</i>.</p>
Snapshot	<p>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.</p> <p>The name of the snapshot is supplied in the Recent Tasks messages for the action.</p> <p>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.</p>
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 6-59. Task ID Dialog Box

Option	Description
OK	To close the dialog box without further action, click OK .
Recent Tasks	To view the status of the job and verify that the job finished, click Recent Tasks .

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

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 list when you click **Administration** in the menu, click **Inventory** , then 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

To ensure that you are submitting the action for the right objects, review the following information.

You first retrieve snapshots based on age, then select the snapshots to delete.

Table 6-60. 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 6-61. 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.
Snapshot Create Time	Date and time when the snapshot was created.
Snapshot Age	Age of the snapshot in days.
Datacenter Name	Name of the data center 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 6-62. 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 .

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

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 list when you click **Administration** in the menu, click **Inventory**, then 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

To ensure that you are submitting the action for the right objects, review the following information. You first retrieve snapshots based on age, then select the snapshots to delete.

Table 6-63. 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.

Table 6-63. Retrieve Snapshots (continued)

Option	Description
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 6-64. 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 data center with which the datastore is associated.
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 6-65. 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 .

The Delete Unused Snapshots action creates a job for the retrieve snapshots action, and a job for the delete snapshots action.

Export Guest User Mapping Action

You can create a template CSV file in any selected directory of vRealize Operations Manager VM and enter credentials for VMs of the selected vCenter Servers.

Where You Run the Action

For the supported objects and object levels, this action is available in the following location in vRealize Operations Manager .

- In the **Administration** tab, select the relevant vCenter Server from Cloud Accounts. Click the vertical ellipsis and select **Object Details**. This action is available from the **Actions** menu below the top menu.

Note This action is deprecated and will be removed in the next release.

Action Menu Items

- 1 Enter the **Export CSV Path** and click **Begin Action**.

You can view the status of the action under **History > Recent Tasks**.

- 2 After the action is successful, log in to vRealize Operations Manager VM with any remote session tool and edit the exported guest user mappings CSV template.
- 3 Encrypt the file with the gpg tool available on vRealize Operations Manager VM by running the command:

```
# cd <guestmappings file parent dir>

# gpg --symmetric <guestmappings filename>
```

- a Enter the same password that you entered in the **Guest User Mapping CSV** field when configuring the Service Discovery adapter.

The gpg tool creates a gpg file which is the encrypted version of the plain-text comma-separated-values file next to the CSV file.

Apply Guest User Mapping Action

You can apply the guest user mappings that you have configured on the vCenter Server.

Where You Run the Action

For the supported objects and object levels, this action is available in the following location in vRealize Operations Manager .

- In the **Administration** tab, select the relevant vCenter Server from Cloud Accounts. Click the vertical ellipsis and select **Object Details**. This action is available from the **Actions** menu below the top menu.

Note This action is deprecated and will be removed in the next release.

Prerequisites

Ensure that the Export Guest User Mapping action is performed successfully.

Action Menu Items

- 1 Enter the **Encrypted CSV Path** (gpg file) and the **Status CSV Path**.
- 2 Select the **Overwrite** check box to overwrite the already configured guest user mapping.
- 3 Click **Begin Action**.

You can view the status of the action under **History > Recent Tasks**.

Note If the user mapping for a VM is not successful, review the CSV path that you entered in the **Status CSV Path** field.

Clear Guest User Mapping Action

You can clear the guest user mapping by specifying an encrypted gpg file.

Where You Run the Action

For the supported objects and object levels, this action is available in the following location in vRealize Operations Manager .

- In the **Administration** tab, select the relevant vCenter Server from Cloud Accounts. Click the vertical ellipsis and select **Object Details**. This action is available from the **Actions** menu below the top menu.

Note This action is deprecated and will be removed in the next release.

Prerequisites

Ensure that the Export Guest User Mapping and the Apply Guest User Mapping actions are performed successfully.

Action Menu Items

- 1 Enter the **Encrypted CSV Path** (gpg file) and the **Status CSV Path**.
- 2 Click **Begin Action**.

You can view the status of the action under **History > Recent Tasks**.

Note If the user mapping for a VM is not successful, review the CSV path that you entered in the **Status CSV Path** field.

Configure Included Services Action

You can extend the set of out-of-the-box discoverable services by adding additional service details.

Where You Run the Action

For the supported objects and object levels, this action is available in the following location in vRealize Operations Manager .

- In the **Administration** tab, select the relevant vCenter Server from **Other Accounts** that have the Service Discovery adapter configured. Click the vertical ellipsis and select **Object Details**. This action is available from the **Actions** menu below the top menu.

Note This action is deprecated and will be removed in the next release.

Action Menu Items

- 1 Add the service details in the format: **<service executable>, <port>, <service name>**. For example, **sshd, 22, SSH Service**.
- 2 Click **Begin Action**.

You can view services under **Home > Manage Applications > Discovered Services**.

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 instance, and configured to run actions. See Configure a vCenter Server Cloud Account section in *vRealize Operations Manager Configuration Guide*.

Verify that your vCenter Adapter is configured to connect to the correct vCenter Server instance, and configured to run actions. See *vRealize Operations Manager Configuration Guide* .

- [Actions Do Not Appear on Object](#)

An action might not appear on an object, such as a host or virtual machine, because vRealize Automation is managing that object.

- [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 specify properly 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 vRealize Automation is managing that object.

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 act on the virtual machines that vRealize Automation is not managing.

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. They are not available because the action framework excludes actions on objects that vRealize Automation manages. You cannot turn on or turn off the exclusion of actions on objects that vRealize Automation manages. 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.

Make actions available on the objects in your data center or custom data center in one of two ways. 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 Perform the action 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 on one or more objects, some of the fields are empty.

Cause

There are two possible causes: 1) the VMware vSphere adapter has not collected the data from the vCenter Server instance that manages the object. 2) 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 specify properly 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 omission 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 five-minute collection cycle.

Solution

- 1 After you create a VMware adapter instance, wait an extra five 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 must 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 6-66. 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 6-66. Task List (continued)

Option	Description
Status	<p>State of the task.</p> <p>Possible states include the following values:</p> <ul style="list-style-type: none"> ■ 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 was not able to start or the script 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 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 investigate this status further, 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	<p>Date and time when the task finished.</p> <p>A completed date does not appear if the task failed or if the maximum timeout is reached.</p>
Automated	Indicates whether the action in the task list is 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.
Event Source	<p>The UUID or the name of the event that triggered the action automatically. When an event is triggered that is associated to the recommendation, it triggers the action without the user intervention.</p> <p>For example, 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.</p> <p>An administrator who has the Automation role has permission to automate actions in the Override Alert / Symptom Definitions area of the policy workspace.</p>

Table 6-66. Task List (continued)

Option	Description
Source Type	Authentication source that the user who started the task used when accessing vRealize Operations Manager .
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	<p>ID generated when the task, which included one or more actions, was started.</p> <p>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.</p> <p>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.</p>

The Associated Objects are the objects on which the selected task ran.

Table 6-67. Associated Objects for Selected Task Details

Option	Description
Object Name	<p>Detailed list of objects that are included in the task selected in the task list.</p> <p>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.</p>
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. If the task does not finish successfully, use the logs to identify problems.

Table 6-68. Messages for Selected Task Details

Severity drop-down menu	Limits the messages based on the Severity value.
Filter (Message)	<p>Limits the message in the list to those that match the entered string.</p> <p>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.</p>
Severity	<p>Message level in the logs.</p> <p>The severity includes the following values:</p> <ul style="list-style-type: none"> ■ Information. Messages added to logs as the task is processed. ■ Error. Messages generated during a task failure.

Table 6-68. Messages for Selected Task Details (continued)

Time	Date and time the entry was added to the log.
Message	<p data-bbox="810 333 1023 354">Text of the log entry.</p> <p data-bbox="810 371 1382 457">Use the information in the message to determine why a task failed, and to begin to troubleshoot and resolve the failure.</p> <p data-bbox="810 474 1414 527">The messages appear with the most recent entry at the top of the list if you do not sort the columns.</p>

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 status of 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. That outcome requires 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](#)

You run the Set CPU Resources or Set Memory Resources action and the task fails with an error appearing in the Recent Tasks messages. The reason might be that you entered a value that is greater than the value that your vCenter Server instance supports.

- [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 shutdown VM action task has a `Failed` status in the Recent Task list.

- [Troubleshooting VMware Tools Not Running for a Shutdown VM Action Status](#)

A Shutdown 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 status of 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 latest 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

To determine whether the action completed successfully, check the state of the target object. If it did not complete, 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 the 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 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. It fails even though Power Off Allowed was selected, the virtual machine is running, and the VMware Tools are installed and running.

Cause

The virtual machine must shut down the guest operating system before it powers off the virtual machine to make the requested changes. The shutdown process waits 120 seconds for a response from the target virtual machine, and fails without changing the virtual machine.

Solution

- 1 To determine if it has jobs running that are delaying the implementation of the action, check the target virtual machine in vCenter Server.
- 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. That outcome requires 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. This change can happen when 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 and the value is changed in vCenter Server. However, the target object is left in a state where you cannot power it on using either actions or the 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 the host does not support 4 CPUs. 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 cannot 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

You run the Set CPU Resources or Set Memory Resources action and the task fails with an error appearing in the Recent Tasks messages. The reason might be that you entered a value that is greater than the value that your vCenter Server instance supports.

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 10000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[10000000000] Mhz
```

The error message for this action is similar to this example.

```
RuntimeException, 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 10000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[10000000000] (MB)
```

The error message for this action is similar to this example.

```
RuntimeException, 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 shutdown 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 shutdown 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 shutdown action is not confirmed.

Solution

- ◆ To determine why the guest operating system did not shut down in the allotted time, check its status in vCenter Server.

Troubleshooting VMware Tools Not Running for a Shutdown VM Action Status

A Shutdown 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 is 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 task 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 To determine if the snapshot was deleted, check its status in vCenter Server .
- 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 6-69. 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 <i>List of vRealize Operations Manager Actions</i> .
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 6-70. 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.

Capacity Optimization for Your Managed Environment

7

Capacity Optimization in vRealize Operations Manager is achieved using powerful integrated functions - capacity overview, workload balancing and optimization, repurposing of underutilized resources, and what-if predictive scenarios - to reach optimal system performance.

Capacity planners must assess whether physical capacity is sufficient to meet current or forecasted demand. With robust capacity planning and optimization, you can manage your production capacity effectively as your organization addresses changing requirements. The objective of strategic capacity optimization is to reach an optimal level where production capabilities meet ongoing demand.

vRealize Operations Manager analytics provide precise tracking, measuring and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery. The system monitors stress thresholds and alerts you before potential issues can affect performance. Multiple pre-set reports are available. You can plan capacity based on historical usage, and run what-if scenarios as your requirements expand.

How Capacity Optimization Works

The Capacity Optimization provides four integrated functions - Overview, Reclaim, Workload Optimization, and What-If Scenarios - that give an overview of the status of all data center activity and trending. You can conduct on-the-spot analysis, including drilling down into further detail on any object to identify possible performance problems or anomalies. You can rebalance and optimize compute resources. The system further identifies underutilized workloads (virtual machines) and calculates the potential cost savings that can accrue when these resources are reclaimed to be deployed more effectively. You can interact with and manipulate data and outcomes based on your requirements.

Use the Capacity Optimization and Reclaim features to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

Workload Optimization provides for moving virtual workloads and their file systems dynamically across datastore clusters within a data center or custom data center. You can potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention triggers an alert and automatically runs an action, a data center performs at optimum.

In addition, the What-If Analysis function- can run scenarios that help determine where additional system resources can be brought online.

Note You may see a data center or cluster labeled as optimized when it has few or no days remaining before CPU, memory, or storage is predicted to run out. That is because these are two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

This chapter includes the following topics:

- [Capacity Analytics](#)
- [Example: Excluding VMs from Reclaim Action](#)
- [What-If Analysis: Modeling Workload, Capacity, or Migration Planning](#)
- [Example: Run a What-If Scenario](#)
- [Example: Import Workload from an Existing VM Scenario](#)
- [Allocation Model](#)
- [Capacity Overview](#)
- [Reclaim](#)
- [Reclamation Settings](#)
- [What-If Analysis - Workload Planning: Traditional](#)
- [What-If Analysis - Infrastructure Planning: Traditional](#)
- [What-If Analysis - Workload Planning: Hyperconverged](#)
- [What-If-Analysis - Infrastructure Planning: Hyperconverged](#)
- [What-If-Analysis - Migration Planning: Public Cloud](#)
- [What-If Analysis - Data Center Comparison](#)
- [Custom Profiles in vRealize Operations Manager](#)
- [Custom Data Centers in vRealize Operations Manager](#)

Capacity Analytics

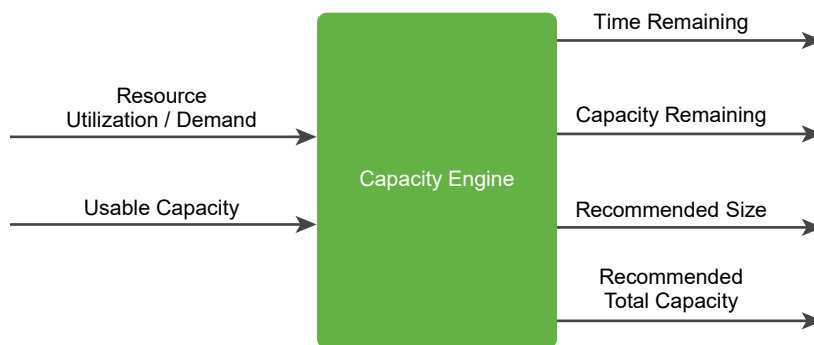
Capacity analytics helps you assess the utilization and capacity remaining in objects across your environment. An evaluation of the historical utilization of resources generates a projection of

the future workload. You can plan for infrastructure procurement or migrations based on the projection and avoid the risk of capacity shortage and high infrastructure costs.

Capacity analytics uses the capacity engine to assess historical trends, which include utilization peaks. The engine chooses an appropriate projection model to predict the future workload. The amount of historical data that is considered depends on the amount of historical utilization data.

Capacity Engine and Calculations

The capacity engine analyzes historical utilization and projects future workload by using real-time predictive capacity analytics, which is based on an industry-standard statistical analysis model of demand behavior. The engine takes the Demand and Usable Capacity metrics as input and generates the output metrics, which are Time Remaining, Capacity Remaining, Recommended Size, and Recommended Total Capacity, as shown in the following figure.



The projection window for the capacity engine is 1 year into the future. The engine consumes data points every 5 minutes to ensure real-time calculation of output metrics.

The capacity engine projects the future workload in a projected utilization range. The range includes an upper bound projection and a lower bound projection. Capacity calculations are based on the time remaining risk level. The engine considers the upper bound projection for a conservative risk level and the mean of the upper bound projection and lower bound projection for an aggressive risk level. For more information about setting risk levels, see, *Capacity Details* in the Configuring Policies chapter of the VMware vRealize Operations Manager Configuration Guide.

The capacity engine calculates the time remaining, capacity remaining, recommended size, and recommended total capacity.

Time Remaining

The number of days remaining till the projected utilization crosses the threshold for the usable capacity. The usable capacity is the total capacity excluding the HA settings.

Capacity Remaining

The largest difference between the usable capacity and the projected utilization between now and 3 days into the future. If the projected utilization is above 100% of the usable capacity, the capacity remaining is 0.

Recommended Size

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. The recommended size excludes HA settings.

If the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization 150 days into the future.

vRealize Operations Manager caps the recommended size that is generated by the capacity engine to keep the recommendations conservative.

- vRealize Operations Manager caps an oversized recommended size at 50% of the currently allocated resources.

For example, a virtual machine that is configured with 8 vCPUs has never used more than 10% CPU historically. Instead of recommending a reclaim of 7 vCPUs, the recommendation is capped to reclaiming 4 vCPUs.

- vRealize Operations Manager caps an undersized recommended size at 100% of the currently allocated resources.

For example, a virtual machine that is configured with 4 vCPUs has been constantly running very hot historically. Instead of recommending the addition of 8 vCPUs, the recommendation is capped at adding 4 vCPUs.


Recommended Total Capacity

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The recommended total capacity includes HA settings.


For example, if the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization including HA values, 150 days into the future.

Note Recommended total capacity is not available for objects.

The following figure shows the capacity calculations for a conservative risk level.



The following figure shows the capacity calculations for an aggressive risk level.



Utilization Peaks

The historical utilization of resources can have peaks, which are periods of maximum utilization. The projection of future workload depends on the types of peaks. According to the frequency of peaks, they can be momentary, sustained, or periodic.

Momentary Peaks

Short-lived peaks that are a one-time occurrence. The peaks are not significant enough to require additional capacity, so they do not impact capacity planning and projection.

Sustained Peaks

Peaks that last for a longer time and impact projections. If a sustained peak is not periodic, the impact on the projection lessens over time because of exponential decay.

Periodic Peaks

Peaks that exhibit cyclical patterns or waves. The peaks can be hourly, daily, weekly, monthly, during the last day of the month, and so on. The capacity engine also detects multiple overlapping cyclical patterns.

Projection Models

The capacity engine uses projection models to generate projections. The engine constantly modifies projections and chooses the model that best fits the pattern of historical data. The projection range predicts the general usage pattern that covers 90% of the future data points. Projection models can be linear or periodic.

Linear Models

Models that have a steadily increasing or decreasing trend. Multiple linear models run in parallel and the capacity engine chooses the best model.

Examples of linear models are linear regression and autoregressive moving average (ARMA).

Periodic Models

Models that discover periodicity of various lengths, such as hours, days, weeks, months, or the last day of the week or month. Periodic models detect square waves that represent batch jobs and handle data streams that contain multiple overlapping periodic patterns. These models ignore random noise.

Examples of periodic models are fast Fourier transforms (FFTs), pulses (edge detection), and wavelets.

Forecast In Trend Views

Forecasts are generated based on the time range specified in the view settings and are forecasted for the number of days specified in the forecast setting. The forecast is generated based on 3 main algorithms. Change-point detection to find sections of the history with significant changes, linear regression to find linear trends, and cyclical analysis to identify periodic patterns.

Historical Data Window

The capacity engine captures historical data over a period of time depending on the historical data window. The historical data window that the engine uses is an exponential decay window.

The exponential decay window is a window of unlimited size in which the capacity engine gives more importance to the most recent data points. Beginning from the projection calculation start point, the engine consumes all the historical data points and weighs them exponentially, based on how far back in time they are.

Example: Excluding VMs from Reclaim Action

In this example, an administrator starts the UI, chooses the Reclaim function on the Quick Start page, and identifies a data center with an excessive number of snapshots. The administrator wants to run the action for reclaiming resources, but chooses to exclude some VMs from the action.

The administrator is reviewing system resources at the start of the shift.

Prerequisites

The administrator must have credentials for operating vRealize Operations Manager and managing vCenter Server objects.

Procedure

- 1 At the Home screen, clicks **Reclaim** in the Optimize Capacity column.

The Reclaim screen appears. In reviewing the status of data centers across the network, the administrator sees that data center DC-Evanston-6 has 3 days of time remaining.

- 2 The administrator clicks the **DC-Evanston-6** graphic.

The data in the lower half of the screen refreshes to display total reclaimable capacity and cost savings potential for recommendations for selected data center DC-Denver-19. (NOTE: Double-clicking the DC-Evanston-6 graphic at this point displays the Object Details page for that data center.)

- 3 At the table, selects **Snapshots** from the header row.

The table refreshes to list clusters with excess snapshots.

- 4 The administrator clicks the **chevron** next to a cluster name on the left in the table.

All the VMs in the cluster are listed.

- 5 The administrator wants to keep snapshots for some VMs in the cluster, so selects two VMs and clicks **EXCLUDE VM(s)**.

A dialog box appears asking for confirmation.

- 6 Clicks **EXCLUDE VM(s)** to confirm.

The excluded VMs disappear from view and the potential cost savings drops.

- 7 Back at the table, with the VMs selected whose snapshots are to be deleted, the administrator clicks **DELETE SNAPSHOT(s)**.

The Delete Snapshots confirmation dialog box appears, showing how many snapshots are to be deleted and the monthly savings in cost and disk space.

- 8 Clicks **DELETE SNAPSHOT(s)** to confirm.

The system deletes the snapshots.

Results

Excessive snapshots are deleted and cost savings are realized.

What to do next

Under Optimize Capacity in the left menu, click **Overview** to display the Capacity Overview screen. Confirm that DC-Evanston-6 now has 15 days of time remaining.

What-If Analysis: Modeling Workload, Capacity, or Migration Planning

Using the what-if tool, you can plan for an increase or decrease in workload or capacity requirements in your virtual infrastructure. To evaluate the demand and supply for capacity on your system objects, and to assess the potential risk to your current capacity, you can create scenarios for adding and removing workloads. You can also determine how much capacity you require to make a migration work. You can run one scenario or group scenarios and run them cumulatively.

Why Create a Scenario

A scenario is a detailed estimation of the resources you must have available in your environment to incorporate upcoming changes. You define scenarios that can potentially add resources to actual data centers. vRealize Operations Manager models the scenario and calculates whether your desired workload can fit in the targeted data center. You can save multiple scenarios for comparison or review.

Where You Find What-If Analysis

From the Home screen, select **What-If Analysis** under **Optimize Capacity** in the left pane. The Overview tab of the What-If analysis page has four panes. Each pane lets you run What-If scenarios to optimize capacity based on workload, physical infrastructure HCI nodes, or migration to the cloud.

How What-If Analysis Works

You can run What-If scenarios to see how much capacity will remain after you add or remove VMs or hosts and add hyperconverged infrastructure (HCI) nodes. Migration planning shows you the capacity and cost information after migrating to cloud based infrastructure.

Scenarios that you save for later are displayed as a list in the **Saved Scenarios** tab. You can run, edit or delete the saved scenarios. You can select more than one compatible scenarios and run them together. For example, you can create a scenario to remove hosts using the **Physical Infrastructure Planning** pane, because your organization has hardware that will soon become obsolete. You can create another scenario to add hosts to your physical infrastructure to account for new hardware that will replace the obsolete ones. You can run both these scenarios together to see the capacity after removing old hardware and adding new hardware.

You can only combine scenarios that pertain to the same object. Use the filters in the **Saved Scenarios** tab to narrow down the list based on scenario name, type, data center, or cluster.

You can select the following combinations of scenarios and run them together:

Workload Planning and Physical Infrastructure Planning

- Add VMs
- Remove VMs
- Add Hosts
- Remove Hosts

The Scenario Summary page displays the results of running one or more saved scenarios. To add or remove saved scenarios and run them again cumulatively, click **Edit** in the **Scenario Summary** page .

Example: Run a What-If Scenario

In this example, an IT administrator at a financial data center must plan for an increase in workloads as tax season approaches. To evaluate whether additional workloads can be added to existing virtual infrastructure, the administrator runs a what-if scenario.

Prerequisites

The administrator must have credentials for operating vRealize Operations Manager and managing vCenter Server objects.

Procedure

- 1 The administrator clicks **Home > Optimize Capacity > What-If Analysis**.

The What-If Analysis screen appears.

- 2 Clicks **Add VMS** in the Workload Planning: Traditional pane.

The Workload Planning: Traditional screen appears.

- 3 Enters Workload Tax 2018 in the **SCENARIO NAME** field, then selects DC-Chicago-16 (vc_10.27.83.19) from the list under **LOCATION - WHERE WOULD YOU LIKE TO ADD YOUR WORKLOAD?**

The field to the right populates with the words, Any cluster. The administrator selects Cluster - Mich2long from the list.

- 4 The administrator clicks the **Configure** radio button.

- 5 For the CPU row, the administrator increments the count to 4. For the Memory row, enters 18. For the Disk Space row, enters 65. Enters 45% in the Expected Utilization column. For number of VMs, enters 20.

The configuration is nearly complete.

6 The administrator clicks **SAVE**

The **Saved Scenarios** screen appears. The data entered on the previous screen appears under Saved Scenarios.

7 The administrator researches the time period for which the workload is needed online.

The administrator identifies the start and end dates.

8 Back at the What-If Analysis screen, the administrator selects Workload Tax 2018 in the list under Saved Scenarios and clicks **EDIT** in the command bar.

The Workload Planning screen appears with the data filled in for the requested scenario.

9 In the **DATE** area, the administrator selects 3/25/18 and 5/30/18 as the start and end dates, respectively, then clicks **RUN SCENARIO**.

The scenario runs and the results appear. To the administrator's surprise, the workload does not fit.

10 At the top right of the screen, the administrator selects a different cluster: Cluster - Mich3long. Then clicks the **RUN SCENARIO** button to the right of the list.

The scenario runs and the results appear. This time the workload fits. It is projected to cost \$84/month to run in the VMware hybrid cloud.

Results

The administrator identifies a location in the virtual infrastructure where the required workload can reside and support the coming increase in production requirements.

What to do next

Assuming this plan is the best of the scenarios the administrator has run, it can be implemented in time to support the added workload. The administrator can monitor the workload performance using the [Using Workload Optimization](#) and [Chapter 7 Capacity Optimization for Your Managed Environment](#) features.

Example: Import Workload from an Existing VM Scenario

In this example, an IT administrator at a data center must plan for an increase in workloads as more staff is hired. To evaluate whether additional workloads can be added to existing virtual infrastructure, the administrator runs a what-if scenario using an actual VM as the workload.

Prerequisites

The administrator must have credentials for operating vRealize Operations Manager and managing vCenter Server objects.

Procedure

1 The administrator clicks **Home > Optimize Capacity > What-If Analysis**.

The What-If Analysis screen appears.

- 2 Clicks **Add VMS** in the Workload Planning: Traditional pane.

The Workload Planning: Traditional screen appears.

- 3 Enters Workload Staff Hire in the **SCENARIO NAME** field, then selects DC-Boston-16 (vc_10.27.83.18) from the list under **LOCATION - WHERE WOULD YOU LIKE TO ADD YOUR WORKLOAD?**

The field to the right populates with the words, Any cluster. The administrator selects Cluster - 1860 from the list.

- 4 The administrator clicks the **Import from existing VM** radio button in the **APPLICATION PROFILE** field, then clicks **SELECT VMs**.

The Select VMs dialog box appears.

- 5 In the column on the left, double-click the name of each VM whose attributes you want use in this scenario. The VM names appear in a **SELECTED** column on the right.

- 6 Click **OK**.

The Workload Planning screen appears. The data entered on the previous screen appears in the **APPLICATION PROFILE** field.

- 7 At the Workload Planning screen, under **APPLICATION PROFILE**, in the **SELECTED VMS** table, enter in the Quantity column the number of copies you want of each VM you selected.

The scenario is almost ready to run.

- 8 In the **DATE** area, the administrator selects 3/25/18 and 6/30/18 as the start and end dates, respectively, then clicks **RUN SCENARIO**

The scenario is successful: the workload will fit. By default, vRealize Operations Manager compares the cost of running the workload on two providers, typically Hybrid Cloud (VMware) and AWS. The corresponding cost details are updated for your private cloud and public cloud providers. The planning scenario also provides a public cloud comparison between Hybrid Cloud and VMware Cloud on AWS. You can see that the monthly cost is displayed for each of the public clouds.

VMware Cloud on AWS	Hybrid Cloud
Shows the number of hosts required on VMare Cloud on AWS for the migration to accommodate the selected workload, considering the minimum purchase of four hosts.	Shows the allocated cost for a month.
The actual utilized capacity of each host, with balanced workload distribution.	Displays the utilization of CPU, memory, and storage. Provides overall requirement of hosts for the given capacity.
Total purchase cost is derived by multiplying the effective monthly purchase cost for each host by the number of required hosts.	
Total Utilized Cost per month is computed based on utilized CPU and RAM, allocated storage, this indicates how well all three resources are being utilized as a fraction of the purchase cost.	

VMware Cloud on AWS	Hybrid Cloud
Required CPU and memory are calculated based on utilization.	
Required storage is calculated based on allocated storage capacity in your private cloud.	
Shows on-demand, one and three-year subscription cost.	
Shows the cost for a selected AWS region and its equivalent resources required for the selected region.	

Results

In the Public Cloud text box, the system displays the monthly cost of running the workload on the VMware Hybrid Cloud versus the AWS Public Cloud.

What to do next

Assuming this plan is the best of the scenarios the administrator has run, it can be implemented in time to support the added workload. The administrator can monitor the workload performance using the [Using Workload Optimization](#) and [Chapter 7 Capacity Optimization for Your Managed Environment](#) features.

Allocation Model

The allocation model determines how much compute, memory, and storage resources are allocated to object types. You define the allocation values by modifying the policy which is applied to the objects. The allocation values, also known as overcommit ratios, affect performance and cost.

The allocation model works alongside the demand model. Unlike the demand model which always affects the capacity calculations, the allocation model can be turned on or off in the policy setting. You can control the ratio by which vRealize Operations Manager overcommits either the CPU, memory, or disk space. By specifying the allocation values in the policy, you can choose whether you want to overcommit your resources or not. Overcommitting helps you measure utilization of resources in a pay-as-you-go model. When you do not overcommit, the utilization of your cluster will never exceed 100%. If your resource utilization is over the allocation ratio that you set, Capacity Remaining becomes zero.

To modify a policy and configure overcommit ratios, see *Policy Allocation Model Element* in *vRealize Operations Manager Configuration Guide*. .

Capacity Overview

Use the Capacity Overview screen to assess workload status and how much capacity is remaining in data centers across your environment.

Where You Find Capacity Overview

In the menu, select home and then click **Overview** under **Optimize Capacity** in the left pane. From the **Quick Start** screen, select **Assess Capacity** in the second-from-left column.

Note Double-click on a data center graphic to display the object details screen for the data center.

How the Capacity Overview Works

The Capacity Optimization and Reclaim features are tightly integrated functions that enable you to assess workload status in data centers across your environment. You can determine time remaining until CPU, memory, or disk space resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the Capacity Overview page, graphical representations of all the data centers and custom data centers in your environment appear. VMware Cloud on AWS data centers has a unique icon to differentiate it from the other data centers.

By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The page refreshes to display the following data:

Time Remaining

Time Remaining specifies which clusters are most constrained and displays the criticality of the cluster.

Optimization Recommendations

vRealize Operations Manager shows you the number of reclaimable VMs and the associated cost savings. Click **View Reclaimable VMs** to navigate to the **Reclaim** page.

Cluster Utilization

Cluster Utilization displays an interactive graph that shows time remaining by component. You can explore the demand percentage over time by CPU, memory, and disk space or by the most constrained component. By default, the data displayed is for the Demand model. If you have configured the Allocation model, then you can also see the CPU, memory, and disk space time remaining model based on the overcommit ratios that you have set in the policy.

Click the **Edit** icon to modify the criticality threshold, risk level, and allocation model. These changes affect the selected cluster's policy. Hence, any change that you make here, affects all the clusters under the same policy.

Set the **Show History** and **Show Forecast** variables to create the slice of time in which you want to see time-remaining data. The vertical axis of the graph shows the total capacity being used by the current amount of CPU, memory, or disk space respectively. The bold, black line across the top of the graph depicts the historical value of usable capacity. The horizontal axis is

the timeline. Vertical lines in the graph are labeled at the bottom of each line. The first vertical dotted line on the left marks the projection calculation start point. The next line is the current date - now. The third vertical marks the date the resource runs out. If a resource has little time remaining, the current date and the date that time runs out may be the same.

vRealize Operations Manager can make recommendations for increasing time remaining based on the data it receives and these recommendations appear at the bottom of the screen. You might see two options: Option 1 shows what you can achieve by reclaiming resources. Option 2 shows the results of adding capacity.

If you choose to reclaim resources, you can run that process immediately by clicking **RECLAIM RESOURCES**. To see the details or choose additional options before running a reclaim action, review the information provided in the **Optimization Recommendations** pane and then click **VIEW RECLAIMABLE VMS** to go to the **Reclaim** page.

Table 7-1. Capacity Optimization Options

Option	Description
Select a datacenter	Select a data center from the carousel across the top of the page. Information about the datacenter is displayed below.
ALL DATACENTERS X	Toggle: click ALL DATACENTERS on the upper right when you want to switch the view to a filtered list of all data centers. Click X to return to a carousel view of data centers.
View:	Filter results to include data centers, custom data centers, or both. This option appears if you select ALL DATACENTERS on the upper right.
Group BY:	Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. This option appears if you select ALL DATACENTERS on the upper right.
Sort by:	Options (Options appear if you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Alarm clock graphic - lists data centers/custom data centers by time remaining. ■ Dollar sign - lists data centers/custom data centers by potential cost savings. ■ Scales graphic - lists data centers/custom data centers by level of optimization.
Select datacenter or ADD NEW CUSTOM DATACENTER	Options (options appear if you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object. ■ Select ADD NEW CUSTOM DATACENTER to display a dialog box that enables you to define a custom data center.

Table 7-1. Capacity Optimization Options (continued)

Option	Description
Time Remaining	<p>Appears when you select a data center or custom data center from the top of the screen.</p> <p>Gives overview of cluster status, including how many are at:</p> <ul style="list-style-type: none"> ■ Critical ■ Medium ■ Normal ■ Unknown <p>"Critical" can indicate a resource contention, imbalance, or other stress condition. Thresholds you set in the policies define what is critical.</p>
Optimization Recommendations	<p>Lists potential cost savings by reclaiming unused resources.</p> <p>Indicates if workloads can be optimized across clusters.</p> <p>VIEW RECLAIMABLE VMS - displays the Reclaim screen, where you can research and run potential VM reclamation actions.</p> <p>VIEW OPTIMIZATION - displays the Workload Optimization screen, where you can optimize workloads based on your policy settings.</p>
Cluster Utilization and Time Remaining	<p>Overall view of cluster health in the selected data center. You can select a cluster from the list to display information about that cluster, or use the options to sort and filter results. The options you select dictate the data displayed in the graph.</p> <p>Sort by:</p> <ul style="list-style-type: none"> ■ Most Constrained: most constrained element ■ CPU (allocation or demand) ■ Memory (allocation or demand) ■ Disk Space (allocation or demand) <hr/> <p>Note Demand model is always on and is the default.</p> <hr/> <p>Filter: search field.</p> <p>Show History for: The period before forecasting begins (does not impact the forecast calculation).</p> <p>Show Forecast For: The forecast period.</p> <p>How is the criticality determined? Displays the criticality threshold you set for this type of object in the Policies Library.</p> <p>Cluster Time Remaining Settings: Click the Edit icon to edit the default policy for the selected cluster. Change the criticality threshold, risk level, allocation model and capacity buffer. Applying these changes affects all objects in the policy. For more information, see <i>Configuring Policies in the VMware vRealize Operations Manager Configuration Guide</i></p>

Table 7-1. Capacity Optimization Options (continued)

Option	Description
Time Remaining graph	Data shows current and trending resource usage and pinpoints when a given cluster is projected to run out of CPU, memory, or disk space based on the allocation or demand model (default).
Recommendations	<p>Option 1: Reclaim Resources.</p> <p>Shows resources that can be reclaimed to increase time remaining for the selected cluster.</p> <p>RECLAIM RESOURCES - displays the Reclaim screen, where you can research and run potential VM reclamation actions.</p> <p>Option 2: Add Capacity.</p> <p>Shows resources that can be added to increase time remaining.</p>

Note You might see that a data center or cluster is labeled optimized when it has few or no days remaining before CPU, memory, or disk space is predicted to run out. The seemingly odd assessment is due to optimization and time remaining being two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

Reclaim

Use the **Reclaim** screen to identify underutilized workloads and reclaim resources from across your environment.

Where You Find Reclaim

From the **Home** screen, select **Reclaim** under **Optimize Capacity** in the left pane. From the **Quick Start** screen, select **Reclaim** in the second-from-left column.

Note Double-click on a data center graphic to display the object details screen for the data center.

How Reclaim Works

The Capacity Optimization and Reclaim features are tightly integrated functions that enable you to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the **Reclaim** page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The area following refreshes to display details about the

selected data center. The **How much you can potentially save** pane reflects potential capacity savings and indicates a possible cost savings once you have reclaimed underused or powered off VMs. The **Total Reclaimable Capacity** pane gives details of the reclaimable percentages for CPU, memory, and disk space.

The table at the bottom of the page provides important information about the VMs that offer the most cost savings. The VMs are listed by **Powered VMs**, **Idle VMs**, **Snapshots**, and **Orphaned Disks**. The highest priority heading is at the far left. You can specify what information is included in your reclaim action. For example, when you click a column heading, the table lists, by data center and then by VM, the allocated and reclaimable CPUs and memory, respectively. Then, for example, you can select the box next to one or more VM names and click the **EXCLUDE VM(S)** button to keep those VMs from being included in any reclaim action. You can also select VMs to resize.

Reclamation Settings

Select the gear icon next to the page heading to customize Reclamation Settings. This affects all data centers. Using the Reclamation Settings, you can exclude, for example all snapshots from being included in the reclaim action - by deselecting the Snapshots check box. Similarly, you can include or exclude powered-off VMs, idle VMs, and orphaned disks. For more information, see [Reclamation Settings](#).

Note To provide read-only access to the Reclamation Settings page for a user, configure the user role in the Access Control page (Roles tab) under **Administration > Access > Access Control**. Select the **Manage Global Settings** permissions under **Administration > Management** in the **Permissions** pane to grant access to modify the Reclamation Settings page. Unselect the **Manage Global Settings** permissions to grant read-only access.

Run a Reclaim Action

Run a reclaim action as follows:

- 1 In the table headings, **Select** the types of VMs to reclaim.
- 2 **Click** the name of a listed cluster to show its VM list.
- 3 **Select** each VM or snapshot you want to reclaim.
- 4 Click **Delete VM(s)** to reclaim their resources.

Table 7-2. Reclaim Options

Option	Description
Select a data center.	Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.
ALL DATACENTERS X	Toggle: click ALL DATACENTERS on the upper right when you want to switch the view to a filtered list of all data centers. Click X to return to a carousel view of data centers.

Table 7-2. Reclaim Options (continued)

Option	Description
View:	Filter results to include data centers, custom data centers, or both. Option appears when you select ALL DATACENTERS on the upper right.
Group BY:	Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select ALL DATACENTERS on the upper right.
Sort by:	Options (Options appear when you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Alarm clock graphic - list data centers/custom data centers by time remaining. ■ Dollar sign - list data centers/custom data centers by potential cost savings. ■ Scales graphic - list data centers/custom data centers by level of optimization.
Select data center or ADD NEW CUSTOM DATACENTER.	Options (Options appear when you select ALL DATACENTERS on the upper right): <ul style="list-style-type: none"> ■ Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. ■ Select ADD NEW CUSTOM DATACENTER to display a dialog box that enables you to define a custom data center.
How much you can potentially save.	Appears when you select a data center or custom data center from the top of the screen. Shows the total calculated potential cost savings when you accept system reclamation recommendations.
Total Reclaimable Capacity	Lists potential cost savings for the selected data center when you reclaim unused resources. Resource: CPU, memory, or disk space Reclaimable Capacity: how much capacity is available to reclaim from idle resources % Reclaimable: percentage of total CPU, memory, or storage you can reclaim.

Table 7-2. Reclaim Options (continued)

Option	Description
Duration older than:	Shows idle or powered off VMs that have been idle or powered off for at least the selected time period: one week, two weeks, or a month.
Table of Potential Cost Savings	<p>Tabular representation of the VMs, Idle VMs, Snapshots, and Orphaned disks in the selected data center from which resources can be reclaimed.</p> <p>Click one of the elements - powered off VMs, idle VMs, and so on - to refresh the table with data for that element. The table lists the relevant clusters. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.</p> <p>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.</p> <p>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.</p> <p>Exclude VM(s): the selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.</p> <p>For powered Off VMs:</p> <ul style="list-style-type: none"> ■ DELETE VM(s): deletes the selected VMs. ■ EXCLUDE VM(s): excludes the selected VMs. <p>For idle VMs:</p> <ul style="list-style-type: none"> ■ DELETE VM(s): deletes the selected VMs. ■ POWER OFF: powers off the selected VMs. ■ EXCLUDE VM(s): excludes the selected VMs. <p>For Snapshots:</p> <ul style="list-style-type: none"> ■ DELETE SNAPSHOT(s): deletes the selected snapshots. ■ EXCLUDE VM(s): excludes the selected snapshot. <p>SHOW/HIDE EXCLUDED VMS: toggle displays or hides the list of VMs you previously excluded.</p> <hr/> <p>Note By default, calculations for reclaimable resources are based on the demand model. But if you turn on the allocation model in the policy settings, the calculations are based on the allocation model.</p> <hr/> <p>For Orphaned Disks:</p> <ul style="list-style-type: none"> ■ EXCLUDE DISK(S): exclude the selected disks in the actionable list. ■ EXPORT ALL: exports the list of orphaned disks into a CSV file. You cannot reclaim orphaned disks from the UI. Instead, export the list into a CSV file and then reclaim the orphaned disks manually. <hr/> <p>Note vRealize Operations Manager reports orphaned VMDKs conservatively. There might be a false positive situation when the used VMDK is reported as orphaned, particularly if the VMDK is located on a datastore which is shared among multiple VCs, while not all the VCs are monitored by vRealize Operations Manager .</p> <p>Check the accuracy of the VMDK reported as an orphaned disk, and then perform a reclamation.</p> <hr/> <p>SHOW/HIDE EXCLUDED DISKS: toggle displays or hides the list of disks you previously excluded. Excluded disks are not listed in the exported CSV file.</p>

Reclamation Settings

Displays information about powered off VMs, idle VMs, snapshots and orphaned disks. This information helps to identify the amount of resources that can be reclaimed and provisioned to other objects in your environment or amount of potential savings that can be done in each month.

The types of VMs are ranked in the order of their importance in a reclamation action. A VM whose attributes match more than one VM type is included with the higher-ranking VM type. Grouping the VMs this way eliminates duplicates during calculations. As an example, powered-off VMs are ranked higher than snapshots, so that a powered-off VM that also has a snapshot appears only in the powered-off VM group.

If you exclude a given type of VM, all VMs matching this type are included with the next lower-ranked group they match. For example, to list all snapshots regardless of whether their corresponding VMs are powered-off or idle, deselect the Powered-off VMs and Idle VMs check boxes.

Further, you can configure how long a given class of VMs must be in the designated state - powered-off, for example, or idle - to be included in the reclamation exercise. You also can choose to hide the cost savings calculation.

Table 7-3. Reclamation Settings

Property	Description
Show Cost Savings	Controls whether to show Cost savings in 'Assess Capacity' and 'Reclaim' pages.
Powered-Off VMs	<p>VMs that have been continuously powered off during the defined period of time.</p> <p>The total storage capacity used is reclaimable. Total storage reclaimable cost is computed by multiplying storage rate with storage utilization. The direct cost of VM is also attributed.</p>
Idle VMs	<p>VMs that have used no more than 100MHz CPU during the defined period of time.</p> <p>Total CPU, memory, and storage capacity allocated to the VMs is reclaimable. Resource level costs are computed by multiplying resource base rate with utilization levels. Direct cost of VM is also attributed.</p>

Table 7-3. Reclamation Settings (continued)

Property	Description
Snapshots	<p>VM snapshots that have existed for the entire defined period of time.</p> <p>Snapshots of a VM use storage space and such storage is reclaimable. The reclaimable cost is computed by multiplying storage rate with reclaimable storage value.</p>
Orphaned Disks	<p>VMDKs on datastores that are not connected to any registered VMs and have not been modified during the defined period of time.</p> <p>Orphaned disks are VMDKs which are associated with a VM which are not in inventory, but still available in a datastore. You can configure the minimum number of days for which VMDKs not related to any existing VM will be reported as orphaned and appear under Orphaned Disks in Reclaim page.</p> <p>Note You can navigate to Global Settings under Administration > Management and change the value of the Orphaned Disks Collection time. At this time that you set, vRealize Operations Manager checks for orphaned VMDKs in vSphere Client instances. The settings for Cost Calculation and Orphaned Disks Collection are interrelated. The default value for Cost Calculation is 9:00 PM, and the default for Orphaned Disks Collection is 8:00 PM. It is recommended to schedule Cost Calculation after Orphaned Disks Collection.</p>

Note If you are unable to make changes in the Reclamation Settings page, your user role in the Access Control page (Roles tab) under **Administration > Access > Access Control** must be modified by an administrator. The **Manage Global Settings** permissions under **Administration > Management** in the **Permissions** pane controls access to the Reclamation Settings page.

What-If Analysis - Workload Planning: Traditional

You define scenarios that can potentially add workloads to actual data centers. vRealize Operations Manager models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center. You can also define scenarios that can potentially remove workloads from data centers. vRealize Operations Manager calculates the time remaining and capacity remaining on the cluster when workloads are removed from the cluster.

Where You Find What-If Analysis - Workload Planning: Traditional

From the Home screen, select **What-If Analysis** under Optimize Capacity in the left pane. From the What-If Analysis screen, click **Add VMs** or **Remove VMs** in the pane titled Workload Planning: Traditional.

How What-If Analysis - Workload Planning: Traditional Works

Capacity Optimization enables you to forecast successfully the impact of adding a workload to an application. By trying various scenarios, you can arrive at an optimum configuration. When you add VMs in the Workload Planning: Traditional pane, you can select the exact data center or custom data center where you want to locate the new workload. You can even pick a specific cluster where the workload is to reside.

In selecting the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage. You have the further option to click Advanced Configuration and specify more precise characteristics for your workload.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the new workload, enter the start and end date for the period when you want the workload to be active. The default is: starting today and ending one year from today. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the vRealize Operations Manager analysis and assessment of your plan.

The system lets you know immediately if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. If you select scenario details, the system displays a graphic depiction of resource use. For each attribute value - vCPU, memory, and storage - the amount by which the workload increases the percentage of total application capacity used is shown against a time line. The graph shows the existing percentage used in blue and the total of existing usage and added usage as a percentage of total capacity in green.

If the proposed workload does not fit, the system announces the outcome and provides the following information:

- How much the added workload reduces the time remaining for the target cluster, for example, from one year to zero.
- The discrepancy between the space available in the target cluster and what the proposed workload requires, for example, 100 GB of memory.
- The cost of the workload on the VMware Hybrid Cloud and on the public cloud.

About Clouds

When you run a scenario in What-If Analysis, you get a recommendation based on cost relative to workload placement on different clouds. This cost-based recommendation varies for different clouds.

Private Cloud and VMware Cloud on AWS costs are computed based on resource usage levels.

Public clouds, AWS, IBM Cloud, Google Cloud, Microsoft Azure, and user-defined cloud costs are dependent on the selected configuration, that is, for the allocated resources. These public cloud instances are selected based on the close proximity rule, with simulated resource allocation values and in some scenarios, the exact configuration match available in the cloud instance list is not available. Due to this issue, these public cloud costs can be inherently higher in comparison.

How What-If Analysis - Remove Workload Works

This feature of Capacity Optimization enables you to forecast successfully the impact of removing a workload. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Workload Planning screen, you can select VMs from the concrete cluster data center or from the customer data center from which you want to remove the existing workload.

While removing workloads, you have two options to define the workload:

- Select existing VMs and use their projected utilization to evaluate the impact of removing workloads.
- Configure the workload manually by specifying the vCPUs, memory, storage, and expected use percentage.

Enter the start and end date for the period during which you want the workload to be removed. By default, the start date is today and the end date is one year from today. The end date is left empty by default. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the vRealize Operations Manager analysis and assessment of your plan.

Table 7-4. What-If Analysis Workload Page Options

Option	Description
Add/Remove VMs	Click Add VMs or Remove VMs to create a scenario for adding or removing workload. When clicked, the command displays the Add Workload or Remove Workload screen.
Scenario Name	In the heading of the Saved Scenarios table. Selecting the check box next to the name selects all scenarios in the list and turns on the dimmed Delete button.
Scenario type	Name of the scenario type. Values are Add Workload, Remove Workload, Add Capacity, Remove Capacity, and Migrate.
<scenario_name>	Name of a saved scenario. Selecting the check box next to a name turns on the dimmed Run Scenario , Edit , and Delete buttons.
All Filters	Use the filter to search for a specific scenario by name or type.
Show Columns	Click the small button on the lower left to display the Show Columns dialog box. You can select up to four columns to display in the table: Scenario Name, Scenario Type, Date Created, and Scenario Start and End Date.

Add or Remove VMs

As part of the What-If workload planning for traditional infrastructure, Workload Planning: Traditional is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. You also have an advanced configuration option that lets you define your configuration more precisely.

Where You Can Add or Remove VMS

At the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional pane.

Table 7-5. Workload Planning: Traditional Add VMs Options

Option	Description
Scenario Name	Name of your scenario
Location	Where do you want to add the workload? Select from the list of existing data centers. You can optionally select the exact cluster where you want the workload to reside.
Application Profile/Configure	Allows you to configure the virtual compute resource, including vCPU, memory, and storage.
Application Profile/Import Import from existing VM	Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload.
Choose Your Workload: <ul style="list-style-type: none"> ■ CPU ■ Memory ■ Disk space 	With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space.
Expected Utilization	Set the projected percentage of total workload capacity you expect to average. Click Advanced Configuration to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning.
Annual Projected Growth	Set the percentage by which you expect your capacity go grow, annually. Click Advanced Configuration to set the percentage growth of CPU, Memory, and Disk individually. For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110. The Annual Projected Growth can be set to 0% if no growth is expected.
Number of VMs (optional)/Quantity	You can optionally select how many VMs to spread the workload across.
Start Date/End Date	Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date.
Run Scenario	Click to run the scenario. The system calculates whether it fits into the location you selected.

Table 7-5. Workload Planning: Traditional Add VMs Options (continued)

Option	Description
Save	Save the scenario.
Cancel	Cancel the scenario.

Table 7-6. Workload Planning: Traditional Remove VMs Options

Option	Description
Scenario Name	Name of your scenario.
Location	From where do you want to remove the workload? Select from the list of existing data centers. You can optionally choose the exact cluster from where you want to remove the workload.
Application Profile/Configure	Allows you to configure the virtual compute resource, including vCPU, memory, and storage. After you have configured the scenario, enter the quantity of custom VMs that you want to remove.
Application Profile/Import Existing VMs	Displays the Select VMs dialog box where you can choose one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to remove from your workload. Note The recommended limit is 100 VMs as a maximum for workload removal.
Application Profile / Custom: Choose your workload ■ CPU ■ Memory ■ Disk space	With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space.
Start Date/End Date	Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank.
Run Scenario	Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload.
Save	Save the scenario.
Cancel	Cancel the scenario.

Select VMs

Use the **Select VMs** dialog box to choose the VMs whose attributes you want to copy or remove for your Workload Planning: Traditional or Workload Planning: Hyperconverged what-if scenarios.

Where You Find Select VMs

From the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional or Workload Planning: Hyperconverged pane. When you have entered a **Scenario Name** and **Location**, click the **Import from existing VM/Existing VMS** radio button, then click **Select VMS**. On the left is a selection box that allows you optionally to choose all VMs. To add a VM to the selected list on the right, double-click on the VM name. Following are the rest of your options:

Select VMs

Option	Description
All Filters	Filter options: VM Name: name of the VM you want. vCenter: all VMs in this vCenter. VM Tag: all VMs with this tag. Custom Group: all VMs in this custom group.
Select (nn).	Select the VMs listed on the current page, from which to import, or remove characteristics.
Select all (nn) VMS	Click to select all the VMs across all the pages, based on the filters you have set. The number of VMs that you can select by clicking this option is limited to 500 VMs.
Selected	List of VMs you selected from RESULTS.
OK	When you have selected the VMs you want, click OK to return to the Add Workload or Remove Workload screen, where your selected VMs are listed.

Under Application Profile, in the Selected VMs table, enter the number of copies of each VM you selected to add or remove in the Quantity column.

Advanced Configuration - Workload

The Advanced Configuration workspace allows you to more precisely define the attributes of the workload you want to use in your what-if analysis.

Where You Find Advanced Configuration

From the What-If Analysis screen, click **Add**. When you have entered a **Scenario Name** and **Location**, click the **Configure** radio button, then click **Advanced Configuration**.

Advanced Configuration Options

Option	Description
Resource Amount	Enter the number of vCPUS, the amount of memory, and the number of storage GBs to include in your scenario configuration.
Expected Utilization	For CPUs, memory, and storage units, respectively, increment the relevant counter to the percentage of total potential usage you expect the resource to use.
Disk space provisioning	Click the radio button for Thin or Thick provisioning.

What-If Analysis - Infrastructure Planning: Traditional

You define scenarios that can potentially add capacity to actual data centers or remove capacity from actual data centers. vRealize Operations Manager models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center.

Where You Find Infrastructure Planning: Traditional

From the Home screen, select **What-If Analysis** under Optimize Capacity in the left pane. Click **Add Hosts** or **Remove Hosts** in the pane titled Infrastructure Planning: Traditional.

How the What-If Analysis for Infrastructure Planning: Traditional Works

Infrastructure Planning for traditional environments enables you to forecast successfully the impact of adding capacity to your environment or removing capacity from your environment. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Infrastructure Planning: Traditional pane, you can choose where you want to locate the additional capacity or from where you can remove the existing capacity.

In selecting the profile while removing capacity, you can select a profile only from server types that exist in your cluster.

In selecting the profile while adding capacity, you have two options:

- Select a server type from a list of commercially available servers. You can select from a list of 1) server types already in your cluster or 2) all server types approved for purchase.
- Configure a custom server manually by specifying CPU attributes, memory, and cost.

When you have set the profile for the new server, enter the number of servers to purchase or remove and the start and end date for the period when you want the scenario to be active. The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. The system can project scenarios ending up to one year from the current date. By default, the starting date is today and the ending date is one year from today.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the vRealize Operations Manager analysis and assessment of your plan.

The system displays immediately the impact on cluster size of the additional or lesser amount of CPU and memory, and shows the total cost of adding or removing the specified capacity. The system also shows whether adding new capacity or removing capacity extends or shrinks the time remaining before CPU or memory runs out.

As well, the system displays a graphic depiction of resource use. For each attribute value - CPU and memory - the amount by which the workload increases or decreases the percentage of total capacity used is shown against a time line.

Add or Remove Hosts

As part of the What-If analysis for physical infrastructure planning for traditional environments, Infrastructure Planning: Traditional pane is what you use to fill in the details of your What-If scenario. You select where to add or remove hosts, use an existing server type, or configure it yourself (when you add capacity), and establish a time frame.

Where You Find Physical Infrastructure

At the What-If Analysis screen, click **Add Hosts** or **Remove Hosts** in the Infrastructure Planning: Traditional pane.

Table 7-7. Add Hosts Options

Option	Description
Scenario Name	Name of your scenario
Location	Where do you want to add capacity? Select from the list of existing data centers, then select the cluster where you want one or more servers to reside.
Server Details	Clicking Select Server displays the Select Server Type dialog box, where you can select a commercial brand server or configure a custom server. Number of Servers to add: increment the Quantity counter up to the number of servers you want.
Start Date/End Date	Select from pop-up calendars the start and end date for the What-If scenario.
Run Scenario	Click to run the scenario. The system calculates the cost of the scenario and determines any new time remaining number.
Save	Save the scenario.
Cancel	Cancel the scenario.

The system displays immediately the impact on cluster size of the additional CPU and memory, and shows the total cost of adding the specified capacity. The system also shows in graphical form whether adding the new capacity extends the time remaining before CPU or memory runs out.

Table 7-8. Remove Hosts Options

Option	Description
Scenario Name	Name of your scenario
Location	From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove one or more servers.
Server Details	Clicking Select Server displays the Select Server Type dialog box, where you can select only the server types that exist in your selected cluster. The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster.
Start Date/End Date	Select from pop-up calendars the start and end date for the What-If scenario. You can select to keep the end date blank.
Run Scenario	Click to run the scenario. The system determines any new time remaining number.

Table 7-8. Remove Hosts Options (continued)

Option	Description
Save	Save the scenario.
Cancel	Cancel the scenario.

The system displays the time remaining and the impact on CPU and memory with reduced capacity. The system also shows in graphical form whether removing capacity decreases the time remaining before CPU or memory runs out.

You can also see that the cost is based on the original purchase cost.

What-If Analysis - Workload Planning: Hyperconverged

You can perform Hyperconverged Infrastructure workload planning by adding or removing VMs to VMware vSAN enabled clusters and running What-If scenarios. vRealize Operations Manager shows you if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. .

Where You Find What-If Analysis - Workload Planning: Hyperconverged

From the menu, select **Home** and **Optimize Capacity > What-If Analysis** in the left pane. From the **What-If Analysis** page, select **Workload Planning: Hyperconverged**. To run a What-If scenario click **Add VMS** or **Remove VMS**.

How What-If Analysis - Workload Planning: Hyperconverged Works

You define scenarios that can potentially add or remove workloads to VMware vSAN environment. The workload scenarios are based on VMs associated with specific storage policy related factors (such as FTT, RAID).

Note When a workload is added based on imported VMs, and the VM is currently in a VMware vSAN-enabled cluster, the VMware vSAN policy settings are not applied and the current VM disk space is taken as is.

Add or Remove VMS

As part of the What-If workload planning for hyperconverged infrastructure, Workload Planning: Hyperconverged is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. The advanced configuration option lets you define your configuration more precisely.

Where You Find Workload Planning

From the menu, select **Home** and **Optimize Capacity > What-If Analysis** in the left pane. Click **Add VMS** or **Remove VMS** in the **Workload Planning: Hyperconverged** pane.

Table 7-9. Workload Planning: Hyperconverged Add Options

Option	Description
Scenario Name	Name of your scenario
Location	Where do you want to add the virtual machines? Select from the list of existing data centers. You can optionally select the exact cluster where you want the virtual machine to reside.
Application Profile/Configure	Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space.
Application Profile/Import Import from existing VM	Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to incorporate as templates into your workload.
Select your workload: <input type="checkbox"/> CPU <input type="checkbox"/> Memory <input type="checkbox"/> Disk space	With the Configure radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space.
Expected Utilization	Set the projected percentage of total workload capacity you expect to average. Click Advanced Configuration to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning.
Annual Projected Growth	Set the percentage by which you expect your capacity to grow, annually. Click Advanced Configuration to set the percentage growth of CPU, Memory, and Disk individually. For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110. The Annual Projected Growth can be set to 0% if no growth is expected.
Number of VMs (optional)/Quantity	You can optionally select how many VMs to spread the workload across.
Additional vSAN configuration	Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup.
Start Date/End Date	Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date.
Run Scenario	Click to run the scenario. The system calculates whether it fits into the location you selected.
Save	Save the scenario.
Cancel	Cancel the scenario.

Table 7-10. Workload Planning: Hyperconverged Remove Options

Option	Description
Scenario Name	Name of your scenario.
Location	From where do you want to remove the VMs? Select from the list of existing data centers. You can optionally select the exact cluster from where you want to remove the workload.
Application Profile/Configure	Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space. After you have configured the scenario, enter the quantity of custom VMs that you want to remove.
Application Profile/Import Existing VMs	Displays the Select VMs dialog box where you can select one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to remove from your workload. Note The recommended limit is 100 VMs as a maximum for workload removal.
Application Profile / Custom: Choose your workload ■ CPU ■ Memory ■ Disk space	With the Configure radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space.
Expected Utilization	Set the projected percentage of total workload capacity you expect to average. Click Advanced Configuration to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning.
Number of VMs (optional)/ Quantity	You can optionally select how many VMs to spread the workload across.
Additional vSAN configuration	Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup.
Start Date/End Date	Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank.
Run Scenario	Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload.
Save	Save the scenario.
Cancel	Cancel the scenario.

What-If-Analysis - Infrastructure Planning: Hyperconverged

You can perform infrastructure planning by adding or removing Hyperconverged Infrastructure (HCI) nodes in vSAN enabled clusters and running What-If scenarios. vRealize Operations Manager displays the cost, time remaining, and capacity remaining for CPU, memory, and disk space in the scenario results.

Where You Find What-If Analysis - Hyperconverged Infrastructure

From the Home screen, select **What-If Analysis** under Optimize Capacity in the left pane. From the What-If Analysis screen, select **Infrastructure Planning: Hyperconverged**. To run a What-If scenario click **Add HCI Nodes** or **Remove HCI Nodes**.

How What-If Analysis - Hyperconverged Infrastructure Works

You can add hyperconverged infrastructure to your VMware vSAN enabled environment evaluate the increase in HCI capacity and cost. You can add up to 64 hosts per vSAN cluster. This number accounts for existing hosts in the cluster. vRealize Operations Manager only lists vSAN and vXRail clusters in the location property. You can select existing server types from these locations and change the number of instances of these servers to add to your scenario.

Note VMC clusters are not supported and will not show up.

Add or Remove HCI Nodes

As part of the what-if analysis for physical infrastructure planning for hyperconverged environments, the Infrastructure Planning: Hyperconverged pane is what you use to fill in the details of your what-if scenario. When you add an HCI node, you can select an existing server type from your vSAN enabled data center and change the number of instances of this server to calculate storage, compute capacity, time remaining, and cost. You can run the Remove HCI Nodes scenario to see the capacity changes after you remove HCI nodes from your data center.

Where You Find Workload Planning

At the **What-If Analysis** page, click **Add HCI Nodes** or **Remove HCI Nodes** in the **Infrastructure Planning: Hyperconverged** pane.

Table 7-11. Add HCI Nodes Options

Option	Description
Scenario Name	Name of your scenario.
Location	Where do you want to add the HCI node? Select from the list of existing data centers. You must also choose the exact cluster where you want the HCI node to reside.
Server Details	Allows you to select an existing server type to calculate capacity, time, and storage remaining based on the number of instances of the server.
Number of servers to add	How many instances of the server do you want to add? Note Only 60 new hosts can be added to the specified vSAN cluster as the maximum allowed is 64 hosts.
Start Date/End Date	Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date.
Run Scenario	Click to run the scenario. The system calculates whether it fits into the location you selected.

Table 7-11. Add HCI Nodes Options (continued)

Option	Description
Save	Save the scenario.
Cancel	Cancel the scenario.

Table 7-12. Remove HCI Nodes Options

Option	Description
Scenario Name	Name of your scenario.
Location	From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove the server(s).
Server Details	Clicking Select Server displays the Select Server Type dialog box, where you can choose only the server types that exist in your selected cluster. The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster.
Start Date/End Date	Select from pop-up calendars the start and end date for the what-if scenario. You can choose to keep the end date blank.
Run Scenario	Click to run the scenario. The system determines any new time remaining number.
Save	Save the scenario.
Cancel	Cancel the scenario.

What-If-Analysis - Migration Planning: Public Cloud

You define scenarios that can potentially migrate workloads to a public cloud or to VMware Cloud on AWS. Use this scenario to determine where to move the workloads. vRealize Operations Manager models the scenario and calculates the cost and capacity to fit your desired workload.

Where You Find What-If Analysis - Migration Planning

From the Home screen, select **What-If Analysis** under Optimize Capacity in the left pane. From the Quick Start screen, select **Plan** in the second-from-left column. Click **Select** in the pane titled Migration Planning.

How What-If Analysis - Migration Planning Works

This feature of Capacity Optimization enables you to forecast successfully the impact of migrating a workload to a public cloud instance such as AWS, IBM Cloud, Microsoft Azure, Google Cloud or to VMware Cloud on AWS. Once you select the Migration Planning screen, choose whether you want to run the scenario against a public cloud or VMware Cloud on AWS. For a public cloud, select the region where you want to migrate the workload. If the public clouds listed out of the box do not suit your needs, you can also define your own public cloud and upload a rate card.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the migrating workload, run the scenario to get the vRealize Operations Manager analysis and assessment of your plan. You can also select up to three public clouds (but not VMware Cloud on AWS) to compare results. Alternatively, you can save the scenario to edit or run later on. A list of saved scenarios is available in the **Saved Scenarios** tab on the What-If analysis page.

For a public cloud target, the system lets you know immediately if the workload proposed for migration fits or does not fit in the suggested location. For example, if you selected AWS and the workload fits, the results list the Amazon Web Services Assessment, with details of the VMware Configuration and the AWS Equivalent. If the proposed workload does not fit, an error message appears: "Unable to identify a matching configuration instance in target location."

If you selected VMware Cloud on AWS for your scenario, the results list the VMware Cloud on AWS Assessment, with details of the VMware configuration. The system also displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription. In addition, the system displays the resource-use-level cost and monthly purchase cost for one-year and three-year subscriptions.

About Clouds

The system might provide a recommendation based on the cost of placing the workload on different clouds. This cost-based recommendation varies for different clouds. You can modify the costs for public clouds by uploading a new rate card.

For VMware Cloud on AWS, the system displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription, plus those same costs for one-year and three-year subscriptions.

Public cloud costs are based on the selected configuration, that is, the allocated resources.

The public instance is selected based on the close proximity rule, with simulated resource allocation values. In some scenarios, an exact configuration match is not available in the list. Due to this lack of availability, the public cost can be inherently higher in comparison.

Migration Planning

As part of the What-If Analysis function, Migrate is the form you use to fill in the details of your what-if scenario. You choose where to migrate the workload, then select the region.

Where You Find Migration Planning

At the What-If Analysis screen, click **SELECT** in the Migrate pane.

When you run a scenario for What If: Migration for Public Clouds (Not VMC), vRealize Operations Manager might suggest the Public Cloud Instance suitable for the Workload Configuration selected by you. vRealize Operations Manager also calculates the cost for that Public Cloud's instance and displays the same.

Table 7-13. Migrate Options

Option	Description
SCENARIO NAME	Name of your scenario
SELECT CLOUDS	<p>Where do you want to migrate the workload?</p> <p>Options:</p> <ul style="list-style-type: none"> ■ AWS ■ VMware Cloud on AWS - You can now select regions for VMware Cloud on AWS. ■ IBM Cloud ■ Microsoft Azure ■ Google Cloud <p>Note The cloud providers added in the Add Cloud Provider page are also included in the list.</p> <p>You can select a maximum of three public clouds at a time for comparison. Hold the Shift key to select more than one public cloud provider. You cannot choose VMware Cloud on AWS with other public clouds for comparison because it has a host-based pricing model, while other clouds are instance-based.</p>
ADD CLOUD PROVIDERS	You can add or edit the cloud providers and also edit the rate card of each individual cloud provider.
APPLICATION PROFILE/Configure	Using the Application Profile you can configure the virtual compute resources, like vCPU, memory, and storage.
Select Your Workload: ■ CPU ■ Memory ■ Disc Space	With the Configure radio button selected, you can size your migrating workload by defining values for vCPU, memory, and storage.
APPLICATION PROFILE/Import from existing VM	<p>Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group.</p> <p>Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload.</p>
Number of VMs (OPTIONAL)/Quantity	You can optionally choose how many VMs to spread the workload across.
RUN SCENARIO	Click to run the scenario. The system calculates whether it fits into the location you chose.

Table 7-13. Migrate Options (continued)

Option	Description
SAVE	SAVE the scenario.
CANCEL	CANCEL the scenario.

VMware Cloud on AWS Assessment - Results

The scenario results are displayed when you run the scenario. For VMware Cloud on AWS Assessment, you can edit the following options.

- **Edit Configuration** - you can edit the change in Reserved Capacity CPU, Reserved Capacity Memory, Fault Tolerance, and RAID Level values and save the values to the original configuration.
- **Change Plan** - you can use the **Choose Plan** option to change your subscription plan, the available options are one-year plan, three-year plan, or Pay-As-You-Go.
- **Edit Discount** - you can use the edit discount option to specify the discount value, the total cost for the subscription is equal to the actual utilization cost minus the discount percentage.

What-If Analysis - Data Center Comparison

You can select virtual machines to determine which of the preferred data centers (along with a specific choice of cluster or default cheapest cluster) are best fit from both cost effectiveness and capacity requirements perspective. The comparison helps you to find the right data center to place the workload from cost and capacity perspective.

Where You Find What-If Analysis – Data Center Comparison

From the Home screen, select **What-If Analysis** under Optimize Capacity in the left pane. From the Quick Start screen, click **Plan** in the second-from-left column. Click **Compare Datacenters** in the pane titled data center comparison.

How What-If Analysis - data center Comparison Works

This feature of Capacity Optimization enables you to compare cost across data centers within the private cloud environment. After you select the Datacenter Comparison screen, choose one or more data centers to compare the cost and run the scenario. vRealize Operations Manager suggests which data center is most cost effective for the selected workload.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying CPU, memory, disk space, expected utilization, and annual projected growth.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for comparing the workload, run the scenario to get the vRealize Operations Manager analysis and assessment of your plan. You can select up to three data centers to compare results. Alternatively, you can save the scenario to edit or run later. A list of saved scenarios is available in the Saved Scenarios tab on the What-If analysis page.

Cost varies from one datacenter to another depending on cost settings, which include cost drivers such as servers, facility, power, labor, license, network, and storage.

The data center comparison feature solves this problem by allowing you to select a data center which suits your requirement, is least expensive, and has adequate capacity.

Datacenter Comparison

As part of the What-If Analysis function, Compare Datacenters is the form you use to fill in the details of your What-If scenario. Use this scenario to compare cost across data centers within the private cloud environment.

Where You Find Compare Datacenters

At the **What-If Analysis** page, click **Compare Datacenters** in the pane titled Datacenter Comparison.

Table 7-14. Compare Datacenter Options

Option	Description
Scenario Name	Name of your scenario.
Select Datacenters	Select the datacenters for which you want to compare the costs.
Application Profile/Configure	Using the Application Profile, you can configure the virtual compute resources, like CPU, memory, disk space, expected utilization, and annual projected growth.
Select Your Workload: ■ CPU ■ Memory ■ Disk Space ■ Expected Utilization ■ Projected Annual Growth	With the Configure radio button selected, you can size your workload by defining values for CPU, memory, disk space, expected utilization, and annual projected growth.
Application Profile/Import from existing VM	Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload.
Number of VMs (OPTIONAL)/Quantity	You can optionally choose how many VMs to spread the workload across.
Date	You can specify the Start Date and End Date to compute the datacenter infrastructure cost for a specific time period.

Table 7-14. Compare Datacenter Options (continued)

Option	Description
Run Scenario	Click to run the scenario. The system calculates the cost of migration and checks whether the selected workload fits into the location you have chosen.
Save	Save the scenario.
Cancel	Cancel the scenario.

Custom Profiles in vRealize Operations Manager

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 remaining and the configuration of that object instance.

To determine how many instances of the object can fit in your environment, use custom profiles with projects and scenarios. Enter the profile numbers or pre-populate the values from specific VMs. 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.

To determine how many instances of the custom profile object you can include on the parent object, you select the parent object and the Capacity tab. The custom profiles appear on the VM remaining section 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.

Where You Find Custom Profiles

In the menu, click **Administration** then **Configuration > Custom Profiles** in the left pane.

Table 7-15. Custom Profiles Options

Option	Description
Toolbar options	In the toolbar click Add Profile to add a custom profile for a specific object type. Click the Vertical Ellipses against a profile to perform the following actions: <ul style="list-style-type: none"> ■ Edit Profile. Modify the selected profile. ■ Delete Profile. Remove the selected profile.
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.
Profile Details tab	Displays the name, description, adapter, object type, and metrics applied to 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 **Add** button. To edit the selected profile, click **Vertical Ellipses** next to the profile and perform an action.

Table 7-16. 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.
Value and Unit	Populate the value and unit for the capacity metrics. You can optionally import the values for an existing VM by clicking the IMPORT FROM EXISTING VM button.

Custom Data Centers in 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 **Capacity** tab.

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 7-17. Custom Datacenters Toolbar and Grid Options

Option	Description
Toolbar options	<p>In the toolbar click Add to add a new custom data center. Click the Vertical Ellipses against a custom data center to perform the following actions:</p> <ul style="list-style-type: none"> ■ Edit. Modify the custom data center. ■ Delete . Remove the custom data center. ■ Clone . Clone the custom data center.
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	<p>Lists the custom data centers in your environment, and displays the health, risk, and efficiency for each one.</p> <p>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.</p>

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 **ADD** button.

To edit a selected custom data center, click the **Vertical Ellipses** to edit, remove or clone.

Table 7-18. 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	<p>Lists the objects in your environment. Select the check box for each object to add to the custom data center.</p> <p>You can add vCenter Server instances, vSphere data centers, vSphere clusters, and ESXi hosts.</p> <p>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.</p>

Metric, Property, and Alert Definitions



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](#).

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 8-1. 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 Analytics Generated Metrics](#).

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Metric Name	Description
CPU Capacity usage	CPU usages as a percent during the interval. Key: cpulcapacity_usagepct_average
CPU CPU contention(%)	<p>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. This number is lower than the highest number experienced by the VM most impacted by CPU contention.</p> <p>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.</p> <p>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.</p> <p>This metric is calculated in the following way: $\text{cpulcapacity_contention} / (200 * \text{summary number_running_vcpus})$</p> <p>Key: cpulcapacity_contentionPct</p>
CPU Demand (%)	<p>This metric shows the amount of CPU resources a virtual machine might use if there were no CPU contention or CPU limit. This metric represents the average active CPU load for the past five minutes.</p> <p>Keep this number below 100% if you set the power management to maximum.</p> <p>This metric is calculated in the following way: $(\text{cpu.demandmhz} / \text{cpu.capacity_provisioned}) * 100$</p> <p>Key: cpuldemandPct</p>
CPU Demand (MHz)	<p>This metric shows the amount of CPU resources a virtual machine might use if there were no CPU contention or CPU limit.</p> <p>Key: cpuldemandmhz</p>
CPU Demand	<p>CPU demand in megahertz.</p> <p>Key: cpuldemand_average</p>
CPU IO wait	<p>IO wait (ms).</p> <p>Key: cpulawait</p>
CPU number of CPU Sockets	<p>Number of CPU sockets.</p> <p>Key: cpulnumpackages</p>
CPU Overall CPU Contention	<p>Overall CPU contention in milliseconds.</p> <p>Key: cpulcapacity_contention</p>
CPU Provisioned Capacity (MHz)	<p>Capacity in MHz of the physical CPU cores.</p> <p>Key: cpulcapacity_provisioned</p>
CPU Provisioned vCPU(s)	<p>Number of provisioned CPU cores.</p> <p>Key: cpulcorecount_provisioned</p>

Metric Name	Description
CPU Reserved Capacity (MHz)	Total CPU capacity reserved by virtual machines. Key: cpulreservedCapacity_average
CPU Usage (MHz)	CPU usages, as measured in megahertz, during the interval. <ul style="list-style-type: none"> ■ 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: $400 / (4 \times 2000) = 0.50$ Key: cpulusagemhz_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: cpulwait
CPU Workload (%)	Percent of workload Key: cpu workload

Memory Metrics

Memory metrics provide information about memory use and allocation.

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

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.

Metric Name	Description
net Packets 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. 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.

Metric Name	Description
disk Total IOPS	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.

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 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 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 Number of Hosts	Total number of hosts. Key: summary total_number_hosts
summary 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 Analytics Generated Metrics](#)
- [Badge Metrics](#)

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Metric Name	Description
Capacity Usage (%)	Percent capacity used. Key: cpulcapacity_usagepct_average
CPU Contention (%)	Percent CPU contention. Key: cpulcapacity_contentionPct
Demand (%)	Percent demand. Key: cpuldemandPct
Demand (MHz)	CPU utilization level based on descendant virtual machines utilization. This Includes reservations, limits, and overhead to run the virtual machines. Key: cpuldemandmhz
Demand	CPU Demand. Key: cpuldemand_average
IO Wait (ms)	IO wait time in milliseconds. Key: cpuliowait
Number of CPU Sockets	Number of CPU sockets. Key: cpulnumpackages
Overall CPU Contention (ms)	Overall CPU contention in milliseconds. Key: cpulcapacity_contention
Provisioned Capacity (MHz)	Provisioned capacity in megahertz. Key: cpulcapacity_provisioned
Provisioned vCPU	Number of provisioned virtual CPU cores. Key: cpulcorecount_provisioned
Reserved Capacity (MHz)	Sum of the reservation properties of the immediate children of the host's root resource pool. Key: cpulreservedCapacity_average
Usage (MHz)	Average CPU use in megahertz. Key: cpulusagemhz_average
Wait (ms)	CPU time spent on the idle state. Key: cpulwait
Overhead	Amount of CPU that is overhead. Key: cpuloverhead_average
Demand without overhead	Value of demand excluding any overhead. Key: cpuldemand_without_overhead
Provisioned Capacity	Provisioned capacity (MHz). Key: cpulvm_capacity_provisioned

Metric Name	Description
Total Capacity (MHz)	Total CPU resources configured on the descendant ESXi hosts. Key: cpulcapacity_provisioned
Usable Capacity (MHz)	The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services. Key: cpulhaTotalCapacity_average

Datastore Metrics

Datastore metrics provide information about the datastore.

Metric Name	Description
Outstanding IO requests	OIO for datastore. Key: datastoreIdemand_oio
Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Read Throughput (KBps)	Amount of data read in the performance interval. Key: datastore read_average
Write Throughput (KBps)	Amount of data written to disk in the performance interval. Key: datastore write_average

Disk Metrics

Disk metrics provide information about disk use.

Metric Name	Description
Total IOPS	Average number of commands issued per second during the collection cycle. Key: disk commandsAveraged_average
Total 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. Key: disk totalLatency_average
Total Throughput (KBps)	Average of the sum of the data read and written for all the disk instances of the host or virtual machine. Key: disk usage_average

Metric Name	Description
Total queued outstanding operations	Sum of queued operations and outstanding operations. Key: disk sum_queued_oio
Max Observed OIO	Max observed IO for a disk. Key: disk max_observed

Disk Space Metrics

Disk space metrics provide information about disk space use.

Metric Name	Description
Total disk space used (KB)	Total disk space used on all datastores visible to this object. Key: disk space total_usage
Total disk space (KB)	Total disk space on all datastores visible to this object. Key: disk space total_capacity
Total provisioned disk space (KB)	Total provisioned disk space on all datastores visible to this object. Key: disk space total_provisioned
Utilization (GB)	Storage space used on the connected vSphere Datastores. Key: disk space total_usage
Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: disk space total_capacity

Memory Metrics

Memory metrics provide information about memory use and allocation.

Metric Name	Description
Contention (%)	Percent host memory contention. Key: mem host_contentionPct
Machine Demand (KB)	Host memory demand in kilobytes. Key: mem host_demand
ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem host_systemUsage
Provisioned Memory (KB)	Provisioned host memory in kilobytes. Key: mem host_provisioned
Reserved Capacity (KB)	Sum of the reservation properties of the immediate children of the host's root resource pool. Key: mem reservedCapacity_average
Usable Memory (KB)	Usable host memory in kilobytes. Key: mem host_usable
Host Usage (KB)	Host memory use in kilobytes. Key: mem host_usage

Metric Name	Description
Usage/Usable (%)	Percent host memory used. Key: mem host_usagePct
Contention (KB)	Host contention in kilobytes. Key: mem host_contention
VM Overhead (KB)	Memory overhead reported by host. Key: mem overhead_average
Utilization (KB)	Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines. Key: mem total_need
Total Capacity (KB)	Total physical memory configured on descendant ESXi hosts. Key: mem host_provisioned
Usable Capacity (KB)	The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services. Key: mem haTotalCapacity_average

Network Metrics

Network metrics provide information about network performance.

Metric Name	Description
Packets Dropped (%)	Percent network packets dropped. Key: net droppedPct
Total Throughput (KBps)	Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine. Key: net usage_average
Packets Received	Number of packets received in the performance interval. Key: net packetsRx_summation
Packets Transmitted	Number of packets transmitted in the performance interval. Key: net packetsTx_summation
Received Packets Dropped	Number of received packets dropped in the performance interval. Key: net droppedRx_summation
Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval. Key: net droppedTx_summation
Data Transmit Rate (KBps)	Average amount of data transmitted per second. Key: net transmitted_average
Data Receive Rate (KBps)	Average amount of data received per second. Key: net received_average

Summary Metrics

Summary metrics provide information about overall performance.

Metric Name	Description
Number of Running Hosts	Number of hosts that are on. Key: summary number_running_hosts
Number of Running VMs	Number of virtual machines that are on. Key: summary number_running_vms
Number of Clusters	Total number of clusters. Key: summary total_number_clusters
Total Number of Datastores	Total number of datastores. Key: summary total_number_datastores
Number of Hosts	Total number of hosts. Key: summary total_number_hosts
Number of VMs	Total number of virtual machines. Key: summary total_number_vms
Maximum Number of VMs	Maximum number of virtual machines. Key: summary max_number_vms
Workload Indicator (%)	Percent workload indicator. Key: summary workload_indicator
Total Number of data centers	Total number of data centers. Key: summary total_number_datacenters
Number of Cores on Powered On Hosts	Number of cores on powered-on hosts. Key: summary number_powered_on_cores
Number VCPUs on Powered on VMs	Number of virtual CPUs on powered-on virtual machines. Key: summary number_running_vcpus
Average Running VM Count per Running Host	Average running virtual machine count per running host. Key: summary avg_vm_density
VC Query Time (ms)	vCenter Server query time in milliseconds. Key: summary vc_query_time
Derived Metrics Computation Time (ms)	Derived metrics computation time in milliseconds. Key: summary derived_metrics_comp_time
Number of objects	Number of objects. Key: summary number_objs
Number of VC Events	Number of vCenter Server events. Key: summary number_vc_events
Number of SMS Metrics	Number of SMS metrics. Key: summary number_sms_metrics
Collector Memory Usage (MB)	Collector memory use in megabytes. Key: summary collector_mem_usage

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Description
Max Observed Number of Outstanding IO Operations	Maximum observed number of outstanding IO operations. Key: datastore maxObserved_OIO
Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore maxObserved_Read
Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore maxObserved_NumberRead
Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore maxObserved_NumberWrite
Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore maxObserved_Write
Max Observed Throughput (KBps)	Max observed rate of network throughput. Key: net maxObserved_KBps
Max Observed Transmitted Throughput (KBps)	Max observed transmitted rate of network throughput. Key: net maxObserved_Tx_KBps
Max Observed Received Throughput (KBps)	Max observed received rate of network throughput. Key: net maxObserved_Rx_KBps

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.

Metrics for ROI Dashboard

Virtual machine metrics provide information about the new metrics added to the ROI dashboard.

Metric Name	Description
Potential Memory Consumed Reclaimable(GB)	This metric displays the sum of all the reclaimable consumed memory for the virtual machine.
Potential CPU Usage Increase (GHz)	This metric displays the potential increase in CPU usage for the virtual machine.
Potential Memory Usage Increase (GB)	This metric displays the potential increase in memory usage for the virtual machine.

Metric Name	Description
Potential Savings	This metric displays the sum of all the potential savings (Idle VMs + Powered off Vms + Snapshot + Orphaned Disks + Oversized VMs).
Potential Cost Increase	This metric displays the potential increase in costs associated with the virtual machine.

Additional Metrics

Twelve new metrics are added to help troubleshoot issues related to virtual machines, POD objects, and Horizon Management Pack objects. You can set the time duration for the metrics collection by enabling or disabling the Near Real Time Monitoring option.

To enable or disable the Near Real Time Monitoring option, click the vertical ellipsis next to the cloud account and select **Edit** option. Select or clear the **Near Real Time Monitoring** option and click **Save**. When you enable the Near Real Time Monitoring option the default time interval is set to 20 seconds and if you disable the Near Real Time Monitoring option the default time interval is set to five minutes .

Note Earlier the average value of the metrics was considered for calculation, now the maximum value is considered for calculation.

Capacity metrics can be calculated for virtual machine objects. See [Capacity Analytics Generated Metrics](#).

Configuration Metrics for Virtual Machines

Configuration metrics provide information about virtual machine configuration.

Metric Name	Description
Config Thin Provisioned Disk	Thin Provisioned Disk. Key: config hardware thin_Enabled
Config Number of CPUs	Number of CPUs for a Virtual Machine. From vRealize Operations Manager 6.7 and onwards, this metric is measured in vCPUs instead of cores. 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 the information about CPU use.

Metric Name	Description
CPU IO Wait (ms)	CPU time spent waiting for IO. Key: cpuliowait
CPU Overall CPU Contention (ms)	The amount of time the CPU cannot run due to contention. Key: cpulcapacity_contention
CPU Reservation Used	CPU Reservation Used. Key: cpu reservation_used
CPU Effective Limit	CPU Effective Limit. Key: cpuleffective_limit
CPU IO Wait (%)	Percentage IO Wait. Key: cpuliowaitPct
CPU Swap wait (%)	Percentage swap waits for CPU. Key: cpulswapwaitPct
CPU Wait (%)	Percentage of the total CPU time spent in wait state. Key: cpu waitPct
CPU System (%)	Percentage CPU time spent on system processes. Key: cpulsystemSummationPct
CPU Capacity entitlement (MHz)	CPU entitlement for the VM after considering all limits. Key: cpulcapacity_entitlement
CPU Capacity Demand Entitlement (%)	Percent capacity demand entitlement. Key: cpulcapacity_demandEntitlementPct
CPU CPU Contention (%)	CPU contention as a percentage of 20-second collection interval. Key: cpulcapacity_contentionPct
CPU Total Capacity	Provisioned CPU capacity in megahertz. Key: cpu vm_capacity_provisioned
CPU Demand (MHz)	Total CPU resources required by the workloads on the virtual machine. Key: cpuldemandmhz
CPU Host demand for aggregation	Host demand for aggregation. Key: cpulhost_demand_for_aggregation
CPU Demand (ms)	The total CPU time that the VM might use if there was no contention. Key: cpuldemand_average
CPU Demand (%)	CPU demand as a percentage of the provisioned capacity. Key: cpuldemandPct
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 when the VM is undersized. Key: cpulusage_average

Metric Name	Description
CPU Usage (MHz)	CPU use in megahertz. Key: cpulusagemhz_average
CPU System (ms)	CPU time spent on system processes. Key: cpulsystem_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. A large ready time 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: cpulreadyPct
CPU Extra (ms)	Extra CPU time in milliseconds. Key: cpulextra_summation
CPU Guaranteed (ms)	CPU time that is guaranteed for the virtual machine. Key: cpulguaranteed_latest
CPU Co-stop (%)	Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpulcostopPct
CPU Latency	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpullatency_average
CPU Max Limited	Time the VM is ready to run, but is not run due to maxing out its CPU limit setting. Key: cpulmaxlimited_summation
CPU Overlap	Time the VM was interrupted to perform system services on behalf of that VM or other VMs. Key: cpuloverlap_summation
CPU Run	Time the VM is scheduled to run. Key: cpulrun_summation
CPU Entitlement Latest	Entitlement Latest. Key: cpulentitlement_latest
CPU Total Capacity (MHz)	Total CPU capacity allocated to the virtual machine. Key: cpulvm_capacity_provisioned
CPU Peak vCPU Ready	The highest CPU Ready among the virtual CPUs. Key: cpulpeak_vcpu_ready
CPU Peak vCPU Usage	The highest CPU Usage among the virtual CPU, compared with the static configured CPU frequency. A constantly high number indicates that one or more of the CPUs have high utilization. Key: cpulpeak_vcpu_usage

Metric Name	Description
CPU 20-second Peak CPU System (%)	The highest CPU system, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak cpu system
CPU 20-second Peak vCPU Co-Stop (%)	The highest CPU Co-Stop among any of the vCPU, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak vcpu co-stop
CPU 20-second Peak vCPU IO-Wait(%)	The highest CPU IO Wait among any of the vCPU, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak vcpu io-wait
CPU 20-second Peak vCPU Overlap (ms)	The highest CPU Overlap among any of the vCPU, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak vcpu overlap
CPU 20-second Peak vCPU Ready (%)	The highest CPU Ready among any of the vCPU, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak vcpu ready
CPU 20-second Peak vCPU Swap Wait (%)	The highest CPU Swap Wait among any of the vCPU, measured as a peak of any 20-second average during the collection interval. Key: cpu 20-second peak vcpu swap wait
CPU vCPU Usage Disparity	The absolute gap between the highest vCPU Usage and the lowest vCPU Usage. Key: cpu vcpu_usage_disparity

CPU Utilization for Resources Metrics for Virtual Machines

CPU utilization for resources metrics provides information about resource CPU use.

Metric Name	Description
rescpu CPU Active (%) (<i>interval</i>)	<p>The average active time (actav) or peak active time (actpk) for the CPU during various intervals.</p> <p>Key:</p> <p>rescpu actav1_latest rescpu actav5_latest rescpu actav15_latest rescpu actpk1_latest rescpu actpk5_latest rescpu actpk15_latest</p>
rescpu CPU Running (%) (<i>interval</i>)	<p>The average runtime (runav) or peak active time (runpk) for the CPU during various intervals.</p> <p>Key:</p> <p>rescpu runav1_latest rescpu runav5_latest rescpu runav15_latest rescpu runpk1_latest rescpu runpk5_latest rescpu runpk15_latest</p>
rescpu CPU Throttled (%) (<i>interval</i>)	<p>Amount of CPU resources over the limit that were refused, average over various intervals.</p> <p>Key:</p> <p>rescpu maxLimited1_latest rescpu maxLimited5_latest rescpu maxLimited15_latest</p>
rescpu Group CPU Sample Count	<p>The sample CPU count.</p> <p>Key: rescpu sampleCount_latest</p>
rescpu Group CPU Sample Period (ms)	<p>The sample period.</p> <p>Key: rescpu samplePeriod_latest</p>

Memory Metrics for Virtual Machines

Memory metrics provide information about memory use and allocation.

Metric Name	Description
Mem Host Active (KB)	<p>Host active memory use in kilobytes.</p> <p>Key: mem host_active</p>
Mem Contention (KB)	<p>Memory contention in kilobytes.</p> <p>Key: mem host_contention</p>
Mem Contention (%)	<p>Percent memory contention.</p> <p>Key: mem host_contentionPct</p>
Mem Guest Configured Memory (KB)	<p>Guest operating system configured memory in kilobytes.</p> <p>Key: mem guest_provisioned</p>

Metric Name	Description
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 Demand for aggregation	Host demand for aggregation. Key: mem host_demand_for_aggregation
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 Host Workload	Host Workload (%). Key: host_workload
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 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

Metric Name	Description
Mem Overhead (KB)	Memory overhead in kilobytes. Key: mem overhead_average
Mem Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem swapiRate_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 compressionRate_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 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 lISwapInRate_average
Mem Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem lISwapOutRate_average
Mem Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem lISwapUsed_average
Mem Overhead Touched	Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM. Key: mem overheadTouched_average
Memory VM Memory Demand (kb)	Key: mem vmMemoryDemand
Memory Consumed (%)	Key: mem consumedPct

Metric Name	Description
Mem Utilization (KB)	Memory used by the virtual machine. Reflects the guest OS memory required for vSphere and certain VMTools versions or for virtual machine consumption. Key: mem vmMemoryDemand
Mem Total Capacity (KB)	Memory resources allocated to powered on virtual machine. Key: mem guest_provisioned
Mem 20-second Peak Contention (%)	The highest Memory Contention, measured as peak of any 20-second average during the collection interval. Key: guest 20-second_peak_contention
Guest Needed Memory	Amount of memory needed for the Guest OS to perform optimally. This memory is considered as a cache for the disk and is a little more than the actual used memory. Key: guest mem.needed_latest
Guest Free Memory	Amount of memory that is not used but is readily available. If the cache is high, a low free memory does not mean that the Guest OS needs more memory. Key: guest mem.free_latest
Guest Physical Usable Memory	Amount of memory available to the Guest OS. Meaning, this amount is close to the amount of configured memory to the VM. Key: guest mem.physUsable_latest
Guest 20-second Peak Disk Queue Length	The highest Disk Queue Length, measured as peak of any 20-second average during the collection interval. Key: guest 20-second_peak_disk_queue_length
Guest 20-second Peak Run Queue	The highest Run Queue, measured as peak of any 20-second average during the collection interval. Key: guest 20-second_peak_run_queue
Guest 20-second Peak CPU Context Switch Rate	The highest CPU Context Switch Rate, measured as peak of any 20-second average during the collection interval. Key: guest 20-second_peak_cpu_context switch rate

Datastore Metrics for Virtual Machines

Datastore metrics provide information about datastore use.

Metric Name	Description
Datastore Total IOPS	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

Metric Name	Description
Datastore Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore io
Datastore Demand	Datastore demand. Key: datastore demand
Datastore Total 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 Total Throughput (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 Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Datastore Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Datastore Read Throughput (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 Throughput (KBps)	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

Disk Metrics for Virtual Machines

Disk metrics provide information about disk use.

Metric Name	Description
Disk Read IOPS	Average number of read commands issued per second during the collection interval. Key: disk numberReadAveraged_average
Disk Write IOPS	Average number of write commands issued per second during the collection interval. Key: disk numberWriteAveraged_average
Disk Total IOPS	Average number of commands issued per second during the collection interval. Key: disk commandsAveraged_average
Disk Total Throughput (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 Throughput KBps)	Amount of data read in the performance interval. Key: disk read_average
Disk Write Throughput (KBps)	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 canceled	The number of disk commands canceled in the performance interval. Key: disk commandsAborted_summation
Disk Highest Latency	Highest latency. Key: disk maxTotalLatency_latest
Disk SCSI Reservation Conflicts	SCSI Reservation Conflicts. Key: disk scsiReservationConflicts_summation

Metric Name	Description
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 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 Total 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 Name	Description
VirtualDisk Usage	Average virtual disk usage as a percentage. Key: virtualDisk usage
VirtualDisk Total Latency	Total latency. Key: virtualDisk totalLatency
VirtualDisk Total IOPS	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 Throughput (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 Name	Description
VirtualDisk Write Throughput (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 canceled 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
Virtual Disk Outstanding IO requests (OIOs)	Key: virtualDisk vDiskOIO
Virtual Disk Used Disk Space (GB)	Key: virtualDisk actualUsage
Virtual Disk Peak Virtual Disk IOPS	The highest disk IO per second among the virtual disks. A constantly high number indicates that one or more virtual disks are sustaining high IOPS. Key: virtualDisk peak_vDisk_iops

Metric Name	Description
Virtual Disk Peak Virtual Disk Read Latency	The highest read latency among the virtual disks. A high number indicates that one or more virtual disks are experiencing poor performance. Key: virtualDisk peak_vDisk_readLatency
Virtual Disk Peak Virtual Disk Write Latency	The highest write latency among the virtual disks. A high number indicates that one or more virtual disks are experiencing poor performance. Key: virtualDisk peak_vDisk_writeLatency
Virtual Disk 20-second Peak Latency (ms)	The highest latency among any of the virtual disk, measured as peak of any 20-second average during the collection interval. Key: virtualDisk 20-second_peak_latency
Virtual Disk Peak Virtual Disk throughput	The highest disk throughput among the virtual disks. Key: virtualDisk peak_vDisk_throughput

Guest File System Metrics for Virtual Machines

Guest file system metrics provide information about guest file system capacity and free space.

The data for these metrics is only displayed when VMware Tools has been installed on the virtual machines. If VMware Tools is not installed, features dependent on these metrics, including capacity planning for virtual machine guest storage, will not be available.

Metric Name	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
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. From vRealize Operations Manager 6.7 and onwards, this metric is measured in GBs. Key: guestfilesystem usage
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

Metric Name	Description
Guest file system Total Guest File System Usage (%)	<p>This metric displays the amount of display space being used out of the total allocated disk space.</p> <p>Use his metric to track if the overall usage is stable, or if it reaches the limits. Do not include VMs with a disk space usage of >95% since this might impact your system.</p> <p>Key: guestfilesystem percentage_total</p>
Guest file system Total Guest File System Usage	<p>Total usage of guest file system.</p> <p>Key: guestfilesystem usage_total</p>
Guest file system Utilization (GB)	<p>Storage space used by the Guest OS file systems. The disk space is available only if VM tools are installed and running. If the VM tools are not installed, the disk space capacity is not applicable.</p> <p>Key: guestfilesystem usage_total</p>
Guest file system Total Capacity (GB)	<p>Storage space used by the Guest OS file systems. The disk space is available only if VM tools are installed and running. If the VM tools are not installed, the disk space capacity is not applicable.</p> <p>Key: guestfilesystem capacity_total</p>

Network Metrics for Virtual Machines

Network metrics provide information about network performance.

Metric Name	Description
Net Total Throughput (KBps)	<p>The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.</p> <p>Key: net usage_average</p>
Net Data Transmit Rate (KBps)	<p>This metric shows the rate of data being sent by the VM per second.</p> <p>Key: net transmitted_average</p>
Net Data Receive Rate (KBps)	<p>This metric shows the rate of data received by the VM per second.</p> <p>Key: net received_average</p>
Net Packets per second	<p>Number of packets transmitted and received per second.</p> <p>Key: net PacketsPerSec</p>
Net Packets Received	<p>Number of packets received in the performance interval.</p> <p>Key: net packetsRx_summation</p>
Net Packets Transmitted	<p>Number of packets transmitted in the performance interval.</p> <p>Key: net packetsTx_summation</p>

Metric Name	Description
Net Transmitted Packets Dropped	This metric shows the number of transmitted packets dropped in the collection interval. Key: net droppedTx_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 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
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 20-second Peak Usage Rate (KBps)	The highest Usage Rate, measured as peak of any 20 second average during the collection interval. Key: net 20-second_peak_usage_rate

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 Name	Description
Sys Powered ON	Powered on virtual machines. 1 if powered on, 0 if powered off, -1 if unknown Key: sys poweredOn
Sys OS Uptime	Total time elapsed, in seconds, since last operating system start. Key: sys osUptime_latest

Power Metrics for Virtual Machines

Power metrics provide information about power use.

Metric Name	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 Name	Description
Diskspace Provisioned Space (GB)	Provisioned space in gigabytes. Key: diskspace provisioned
Diskspace Provisioned Space for VM	Provisioned space for VM. Key: diskspace provisionedSpace
Diskspace Snapshot Space (GB)	Space used by snapshots. Key: diskspace snapshot
Diskspace Virtual machine used (GB)	Space used by virtual machine files in gigabytes. Key: diskspace perDsUsed
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 Name	Description
Storage Total IOPS	Average number of commands issued per second during the collection interval. Key: storage commandsAveraged_average
Storage Contention (%)	Percent contention. Key: storage contention
Storage Read Throughput (KBps)	Read throughput rate in kilobytes per second. Key: storage read_average
Storage Read IOPS	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

Metric Name	Description
Storage Write Throughput (KBps)	Write throughput rate in kilobytes per second. Key: storage write_average
Storage Write IOPS	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 Name	Description
Summary Running	Number of running virtual machines. Key: summary running
Summary Desktop Status	Horizon view desktop status. Key: summary desktop_status
Summary Configuration Type	Indicates the type of virtual machine object based on which you can identify the type of virtual machine. The valid values for the virtual machine object property are: <ul style="list-style-type: none"> ■ default - represents a regular virtual machine ■ template - represents a powered off virtual machine template. ■ srm_placeholder - represents a powered on Site Recovery Manager virtual machine. ■ ft_primary - represents the primary Fault Tolerance virtual machine. ■ ft_secondary - represents the secondary Fault Tolerance virtual machine. Key: summary config type
Summary Guest Operating System Guest OS Full Name	Displays the guest operating system name. Key: summary guest os full name
Summary Oversized Potential Memory	Displays the oversized potential memory. Key: summary oversized potentialMemConsumed
Summary Undersized Potential CPU Usage	Displays the undersized potential CPU used. Key: summary undersized potentialCpuUsage
Summary Undersized Potential Memory	Displays the undersized potential memory used. Key: summary undersized potentialMemUsage
Reclaimable Idle	Boolean flag indicating whether VM is considered as reclaimable because it is in Idle state. Key: summary idle

Metric Name	Description
Reclaimable Powered Off	Boolean flag indicating whether VM is considered as reclaimable because it is in powered off state. Key: summary poweredOff
Reclaimable Snapshot Space (GB)	Reclaimable snapshot space. Key: summary snapshotSpace

Cost Metrics for Virtual Machines

Cost metrics provide information about the cost.

Metric Name	Description
Monthly OS Labor Cost	Monthly operating system labor cost of the virtual machine. Key: cost osLaborTotalCost
Monthly Projected Total Cost	Virtual machine cost projected for full month. Key: Cost monthlyProjectedCost
Monthly VI Labor Cost	Monthly virtual infrastructure labor cost of the virtual machine. Key: cost viLaborTotalCost
MTD Compute Total Cost	Total compute cost (including CPU and memory) of the virtual machine. Key: cost compTotalCost
MTD CPU Cost	Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost. Key: cost cpuCost
MTD Monthly Cost	Month to date direct cost (comprising of OS labor, VI labor and any windows desktop instance license) of the virtual machine. It also comprises of the additional and application cost of the virtual machine. Key: cost vmDirectCost
MTD Memory Cost	Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost. Key: cost memoryCost
MTD Storage Cost	Month to date storage cost of the virtual machine. Key: cost storageCost
MTD Total Cost	Month to date total compute cost (including CPU and memory) of the virtual machine. Key: cost monthlyTotalCost
Potential Savings	Reclaimable cost of VM for being either idle, powered-off, or having snapshots. Key: cost reclaimableCost

Metric Name	Description
Cost Allocation MTD VM CPU Cost (Currency)	Month to Date Virtual Machine CPU Cost computed based on resource overcommit ratio set for its parent cluster in policy. cost allocation allocationBasedCpuMTDCost
Cost Allocation MTD VM Memory Cost (Currency)	Month to Date Virtual Machine CPU Memory cost computed based on resource overcommit ratio set for its parent cluster in policy. cost allocation allocationBasedMemoryMTDCost
Cost Allocation MTD VM Storage Cost (Currency)	Month to Date Virtual Machine CPU Storage cost computed based on resource overcommit ratio set for its parent cluster (or datastore cluster) in policy. cost allocation allocationBasedStorageMTDCost
Cost Allocation MTD VM Total Cost (Currency)	Month to Date Virtual Machine Total Cost is the summation of the CPU Cost, Memory Cost, Storage Cost and Direct Cost, based on overcommit ratios set in policy for the parent cluster or datastore cluster. cost allocation allocationBasedTotalCost
Cost Effective Daily Cpu Cost (Currency)	Daily CPU cost of the selected virtual machine.
Cost Effective Daily Memory Cost (Currency)	Daily Memory cost of the selected virtual machine.
Cost Effective Daily Storage Cost (Currency)	Daily Storage cost of the selected virtual machine.
Cost Daily Additional Cost	Daily Additional cost of the selected virtual machine.
Cost Effective Daily Cost (Currency)	Effective Daily cost is the sum of effective daily CPU cost + effective daily memory cost + effective daily storage cost + daily additional cost.
Cost Effective MTD Cost (Currency)	Effective MTD cost is the sum of effective daily CPU cost from beginning of month until now + effective daily memory cost from beginning of month until now + effective daily storage cost from beginning of month until now + daily additional cost from beginning of month until now.

Virtual Hardware Metrics for Virtual Machines

Metric Name	Description
Configuration Hardware Number of CPU cores per socket	This metric displays the number of CPU cores per socket.
Configuration Hardware Number of virtual CPUs	This metric displays the number of CPUs in the virtual machine.
Configuration Hardware Number of virtual sockets:	This metric displays the number of virtual sockets in the virtual machine.
Configuration Hardware Memory:	This metric displays the memory used in the virtual machine.

Metric Name	Description
Configuration CPU Resource Allocation Limit	This metric displays the resource allocation limit of the virtual machine.
Configuration CPU Resource Allocation Reservation	This metric displays the reserved resources for the virtual machine.
Configuration CPU Resource Allocation Shares	This metric displays the shared resources for the virtual machine.
Summary Guest Operating System Tools Version	This metric displays the tools version of the guest operating system.
Summary Guest Operating System Tools Version Status	This metric displays the status of the tools in the guest operating system.
Summary Guest Operating System Tools Running Status	This metric displays whether the tools are functional in the guest operating system.
Guest File System:/boot Partition Capacity (GB)	This metric displays the boot partition capacity in the guest file system.
Guest File System:/boot Partition Utilization (%)	This metric displays the boot partition usage percentage in the guest file system.
Guest File System:/boot Partition Utilization (GB)	This metric displays the boot partition used in the guest file system.
Virtual Disk Configured	This metric displays the disk space of the configured virtual disk.
Virtual Disk Label	This metric displays the disk label of the configured virtual disk.
Disk Space Snapshot Space	This metric displays the snap shot details of the virtual machine.
Network IP Address	This metric displays the IP address of the virtual machine.
Network MAC Address	This metric displays the MAC address of the virtual machine.

Disabled Instanced Metrics

The instance metrics created for the following metrics are disabled in this version of vRealize Operations Manager . This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

Metric Name
Configuration Hardware Number of virtual CPUs
CPU Ready (%)
CPU Usage (MHz)
Net Broadcast Packets Transmitted
Net Data Transmit Rate (KBps)

Metric Name
Net Data Receive Rate (KBps)
Net Multicast Packets Transmitted
Net Packets Dropped
Net Packets Dropped (%)
Net pnicByteRx_average
Net pnicByteTx_average
Net Transmitted Packets Dropped
Net Usage Rate (KBps)
VirtualDisk Read IOPS
VirtualDisk Read Latency (ms)
VirtualDisk Read Throughput (KBps)
VirtualDisk Total IOPS
VirtualDisk Total Latency
VirtualDisk Total Throughput (KBps)
Virtual Disk Used Disk Space (GB)
VirtualDisk Write IOPS
VirtualDisk Write Latency (ms)
VirtualDisk Write Throughput (KBps)
Datastore Outstanding IO requests
Datastore Read IOPS
Datastore Read Latency (ms)
Datastore Read Throughput (KBps)
Datastore Total IOPS
Datastore Total Latency (ms)
Datastore Total Throughput (KBps)
Datastore Write IOPS
Datastore Write Latency (ms)
Datastore Write Throughput (KBps)
Disk Total IOPS

Metric Name
Disk Total Throughput (KBps)
Disk Read Throughput KBps)
Disk Write Throughput (KBps)
Diskspace Access Time (ms)
Diskspace Virtual machine used (GB)

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Description
CPU 50% of Recommended number of vCPUs to Remove	This metric is superseded by the capacity engine. cpu numberToRemove50Pct
CPU Capacity entitlement (mhz)	cpu capacity_entitlement
CPU Co-stop (msec)	Use the Co-Stop (%) metric instead of this metric. cpu costop_summation
CPU Demand Over Capacity (mhz)	cpu demandOverCapacity
CPU Demand Over Limit (mhz)	Use Contention (%) metric instead of this metric. cpu demandOverLimit
CPU Dynamic entitlement	cpu dynamic_entitlement
CPU Estimated entitlement	cpu estimated_entitlement
CPU Idle (%)	cpu idlePct
CPU Idle (msec)	cpu idle_summation
CPU IO Wait (msec)	cpu iowait
CPU Normalized Co-stop (%)	Use the Co-Stop (%) metric instead of this metric. cpu perCpuCoStopPct
CPU Provisioned vCPU(s) (Cores)	cpu corecount_provisioned
CPU Ready (msec)	Choose the Use Ready (%) metric instead of this metric. cpu ready_summation
CPU Recommended Size Reduction (%)	cpu sizePctReduction
CPU Swap Wait (msec)	cpu swapwait_summation

Metric Name	Description
CPU Total Wait (msec)	cpu wait
CPU Used (msec)	cpu used_summation
CPU Wait (msec)	cpu wait_summation
Datastore I/O Max Observed Number of Outstanding IO Operations	datastore maxObserved_OIO
Datastore I/O Max Observed Read Rate (kbps)	datastore maxObserved_Read
Datastore I/O Max Observed Reads per second	datastore maxObserved_NumberRead
Datastore I/O Max Observed Write Rate (kbps)	datastore maxObserved_Write
Datastore I/O Max Observed Writes per second	datastore maxObserved_NumberWrite
Disk Space Not Shared (gb)	diskspace notshared
Disk Space Number of Virtual Disks	diskspace numvmdisk
Disk Space Shared Used (gb)	diskspace shared
Disk Space Total disk space used (gb)	diskspace total_usage
Disk Space Total disk space (gb)	diskspace total_capacity
Disk Space Virtual Disk Used (gb)	diskspace diskused
Guest File System stats Total Guest File System Free (gab)	guestfilesystem freespace_total
Guest Active File Cache Memory (kb)	guest mem.activeFileCache_latest
Guest Context Swap Rate per second	guest contextSwapRate_latest
Guest Huge Page Size (kb)	guest hugePage.size_latest
Guest Page Out Rate per second	guest page.outRate_latest
Guest Total Huge Pages	guest hugePage.total_latest
Memory 50% of Reclaimable Memory Capacity (gb)	This metric is superseded by the capacity engine. mem wasteValue50PctInGB
Memory Balloon (kb)	mem vmmemctl_average
Memory Demand Over Capacity	mem demandOverCapacity
Memory Demand Over Limit	mem demandOverLimit
Memory Granted (kb)	mem granted_average
Memory Guest Active (kb)	mem active_average
Memory Guest Dynamic Entitlement (kb)	mem guest_dynamic_entitlement
Memory Guest Workload (%)	mem guest_workload

Metric Name	Description
Memory Host Demand with Reservation (kb)	mem host_demand_reservation
Memory Host Dynamic Entitlement (kb)	mem host_dynamic_entitlement
Memory Host Usage (kb)	mem host_usage
Memory Host Workload (%)	mem host_workload
Memory Latency (%)	Use the Memory Contention (%) metric instead of this metric. mem latency_average
Memory Recommended Size Reduction (%)	mem sizePctReduction
Memory Shared (kb)	mem shared_average
Memory Swap Out Rate (kbps)	mem swapoutRate_average
Memory Usage (%)	mem usage_average
Memory Estimated entitlement	mem estimated_entitlement
Network I/O Data Receive Demand Rate (kbps)	net receive_demand_average
Network I/O Data Transmit Demand Rate (kbps)	net transmit_demand_average
Network I/O VM to Host Data Receive Rate (kbps)	net host_received_average
Network I/O VM to Host Data Transmit Rate (kbps)	net host_transmitted_average
Network I/O VM to Host Max Observed Received Throughput (kbps)	net host_maxObserved_Rx_KBps
Network I/O VM to Host Max Observed Throughput (kbps)	net host_maxObserved_KBps
Network I/O VM to Host Max Observed Transmitted Throughput (kbps)	net host_maxObserved_Tx_KBps
Network I/O VM to Host Usage Rate (kbps)	net host_usage_average
Network bytesRx (kbps)	net bytesRx_average
Network bytesTx (kbps)	net bytesTx_average
Network Demand (%)	Use absolute numbers instead of this metric. net demand
Network I/O Usage Capacity	net usage_capacity
Network Max Observed Received Throughput (kbps)	net maxObserved_Rx_KBps
Network Max Observed Throughput (kbps)	net maxObserved_KBps
Network Max Observed Transmitted Throughput (kbps)	net maxObserved_Tx_KBps
Network Packets Received per second	net packetsRxPerSec
Network Packets Transmitted per second	net packetsTxPerSec

Metric Name	Description
Network Received Packets Dropped	net droppedRx_summation
Storage Demand (kbps)	storage demandKBps
Storage Read Latency (msec)	storage totalReadLatency_average
Storage Write Latency (msec)	storage totalWriteLatency_average
Summary CPU Shares	summary cpu_shares
Summary Memory Shares	summary mem_shares
Summary Number of Datastores	summary number_datastore
Summary Number of Networks	summary number_network
Summary Workload Indicator	summary workload_indicator
System Build Number	sys build
System Heartbeat	sys heartbeat_summation
System Product String	sys productString
System Uptime (sec)	sys uptime_latest
System vMotion Enabled	vMotion should be enabled for all. It is not necessary to track all VMs every five minutes. sys vmotionEnabled

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 Analytics Generated Metrics](#).

Configuration Metrics for Host Systems

Configuration metrics provide information about host system configuration.

Metric Name	Description
Configuration Hyperthreading Active	Displays the hyperthreading status of the host. Key: configuration hypwerthreading active
Configuration Hyperthreading Available	Displays whether the hyperthreading option is available for this host. Key: configuration hypwerthreading available
Configuration Storage Device Multipath Info Total number of Active Path	Displays the amount of active path information for the storage device Key: configuration storagedevice multipathinfo total numberofActive path

Metric Name	Description
Configuration Storage Device Total number of path	Displays the total number of path for the storage device. Key: configuration storagedevice total number of path
Configuration Failover Hosts	Failover Hosts. Key: configuration dasConfig admissionControlPolicy failoverHost

Hardware Metrics for Host Systems

Hardware metrics provide information about host system hardware.

Metric Name	Description
Hardware Number of CPUs	Number of CPUs for a host. Key: hardware cpuinfo num_CpuCores
Hardware ServiceTag	Displays the service tag of the host system. Key: hardware servicetag

CPU Usage Metrics for Host Systems

CPU usage metrics provide information about CPU use.

Metric Name	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 (%)	<p>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).</p> <p>Use this metric to verify if the host is able to serve all of its VMs well.</p> <p>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.</p> <p>Absolute means that the real value is high and should be checked. 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 consideration. This is because the VM is not running at full speed.</p> <p>Key: cpu capacity_contentionPct</p>

Metric Name	Description
CPU Demand (%)	<p>This metric shows the percentage of CPU resources all the VMs would use if there was no CPU contention or any CPU limits set.</p> <p>It represents the average active CPU load for the past five minutes.</p> <p>Keep the number of this metric below 100% if you set Power Management to Maximum.</p> <p>Key: cpudemandPct</p>
CPU Demand (MHz)	<p>CPU demand in megahertz. CPU utilization level based on descendant Virtual Machines utilization. Includes limits and overhead to run Virtual Machines, but not reservations.</p> <p>Key: cpudemandmhz</p>
CPU IO Wait (ms)	<p>IO wait time in milliseconds.</p> <p>Key: cpuliowait</p>
CPU Number of CPU Sockets	<p>Number of CPU sockets.</p> <p>Key: cpunumpackages</p>
CPU Overall CPU Contention (ms)	<p>Overall CPU contention in milliseconds.</p> <p>Key: cpulcapacity_contention</p>
CPU Provisioned Capacity (MHz)	<p>Capacity in MHz of the physical CPU cores.</p> <p>Key: cpulcapacity_provisioned</p>
CPU Provisioned virtual CPUs	<p>Provisioned virtual CPUs.</p> <p>Key: cpulcorecount_provisioned</p>
CPU Total Wait	<p>CPU time spent in idle state.</p> <p>Key: cpulwait</p>
CPU Demand	<p>CPU demand.</p> <p>Key: cpuldemand_average</p>
CPU Usage (MHz)	<p>CPU use in megahertz.</p> <p>Key: cpulusagemhz_average</p>
CPU Reserved Capacity (MHz)	<p>The sum of the reservation properties of the (immediate) children of the host's root resource pool.</p> <p>Key: cpulreservedCapacity_average</p>
CPU Total Capacity (MHz)	<p>Total CPU capacity in megahertz. Amount of CPU resources configured on the ESXi hosts.</p> <p>Key: cpulcapacity_provisioned</p>
CPU Overhead (KB)	<p>Amount of CPU overhead.</p> <p>Key: cpuloverhead_average</p>
CPU Demand without overhead	<p>Value of demand excluding any overhead.</p> <p>Key: cpuldemand_without_overhead</p>
CPU Core Utilization (%)	<p>Percent core utilization.</p> <p>Key: cpulcoreUtilization_average</p>

Metric Name	Description
CPU Utilization(%)	Percent CPU utilization. Key: cpulutilization_average
CPU Core Utilization (%)	Core Utilization. Key: cpulcoreUtilization_average
CPU Utilization (%)	Utilization. Key: cpulutilization_average
CPU Co-stop (ms)	Time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpulcostop_summation
CPU Latency (%)	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpullatency_average
CPU Ready (ms)	Time spent in ready state. Key: cpulready_summation
CPU Run (ms)	Time the virtual machine is scheduled to run. Key: cpulrun_summation
CPU Swap wait (ms)	Amount of time waiting for swap space. Key: cpulswapwait_summation
CPU Wait (ms)	Total CPU time spent in wait state. Key: cpulwait_summation
CPU Provisioned Capacity	Provisioned capacity (MHz). Key: cpulvm_capacity_provisioned
CPU Active Host Load For Balance (Long Term)	Active Host Load For Balance (Long Term). Key: cpulacvmWorkloadDisparityPcttive_longterm_load
CPU Active Host Load For Balance (Short Term)	Active Host Load For Balance (Short Term). Key: cpulactive_shortterm_load
CPU CPU Model	Displays the host's CPU model. Key: cpulcpu model
CPU Peak CPU Core Usage	The highest CPU Usage among the CPU cores. A constantly high number indicates that one or more physical cores have high utilization. Key: cpulpeak_cpu_core_usage

CPU Utilization for Resources Metrics for Host Systems

CPU utilization for resources metrics provide information about CPU activity.

Metric Name	Description
Rescpu CPU Active (%) (<i>interval</i>)	<p>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.</p> <p>Key:</p> <p>rescpu actav1_latest rescpu actav5_latest rescpu actav15_latest rescpu actpk1_latest rescpu actpk5_latest rescpu actpk15_latest</p>
Rescpu CPU Running (%) (<i>interval</i>)	<p>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.</p> <p>Key:</p> <p>rescpu runav1_latest rescpu runav5_latest rescpu runav15_latest rescpu runpk1_latest rescpu runpk5_latest rescpu runpk15_latest</p>
Rescpu CPU Throttled (%) (<i>interval</i>)	<p>Scheduling limit over the past minute, past five minutes, and past 15 minutes.</p> <p>Key:</p> <p>rescpu maxLimited1_latest rescpu maxLimited5_latest rescpu maxLimited15_latest</p>
Rescpu Group CPU Sample Count	<p>Group CPU sample count.</p> <p>Key: rescpu sampleCount_latest</p>
Rescpu Group CPU Sample Period (ms)	<p>Group CPU sample period in milliseconds.</p> <p>Key: rescpu samplePeriod_latest</p>

Datastore Metrics for Host Systems

Datastore metrics provide information about datastore use.

Metric Name	Description
Datastore Outstanding IO requests	<p>OIO for datastore.</p> <p>Key: datastore demand_oio</p>
Datastore Commands Averaged	<p>Average number of commands issued per second during the collection interval.</p> <p>Key: datastore commandsAveraged_average</p>
Datastore Number of Outstanding IO Operations	<p>Number of outstanding IO operations.</p> <p>Key: datastore oio</p>

Metric Name	Description
Datastore Total 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 Total Throughput (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
Datastore Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Datastore Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Datastore Read Throughput (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 Throughput (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

Metric Name	Description
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
Datastore Data Written	Data Written. Key: datastore datastoreWriteBytes_latest
Datastore Data Write Rate	Data Write Rate. Key: datastore datastoreWriteIops_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 VM Disk I/O Workload Disparity	Percentage Disk I/O workload disparity among the VMs on the Host. Key: datastore vmWorkloadDisparityPc
Datastore Peak Datastore Read Latency	The highest read latency among the datastores. A high number indicates that one or more datastores are experiencing poor performance. Key: datastore peak_datastore_readLatency
Datastore Peak Datastore Write Latency	The highest write latency among the datastores. A high number indicates that one or more datastores are experiencing poor performance. Key: datastore peak_datastore_writeLatency

Disk Metrics for Host Systems

Disk metrics provide information about disk use.

Metric Name	Description
Disk Total Throughput (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

Metric Name	Description
Disk Total IOPS	Average number of commands issued per second during the collection interval. Key: disk commandsAveraged_average
Disk Total 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 Read IOPS	Average number of read commands issued per second during the collection interval. Key: disk numberReadAveraged_average
Disk Write IOPS	Average number of write commands issued per second during the collection interval. Key: disk numberWriteAveraged_average
Disk Read Throughput (KBps)	Amount of data read in the performance interval. Key: disk read_average
Disk Write Throughput (KBps)	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 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 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 Physical Device Latency (ms)	The average time taken to complete a command from the physical device. Key: disk deviceLatency_average
Disk Kernel Latency (ms)	The average time spent in ESX Server VMKernel per command. Key: disk kernelLatency_average
Disk Queue 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

Metric Name	Description
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 Name	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

Metric Name	Description
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
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 swpin_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

Metric Name	Description
Mem Balloon (KB)	<p>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.</p> <p>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.</p> <p>Key: mem vmemctl_average</p>
Mem Zero (KB)	<p>Amount of memory that is all zero.</p> <p>Key: mem zero_average</p>
Mem State (0-3)	<p>Overall state of the memory. The value is an integer between 0 (high) and 3 (low).</p> <p>Key: mem state_latest</p>
Mem Usage (KB)	<p>Host memory use in kilobytes.</p> <p>Key: mem host_usage</p>
Mem Usage (%)	<p>Memory currently in use as a percentage of total available memory.</p> <p>Key: mem usage_average</p>
Mem Swap In Rate (KBps)	<p>Rate at which memory is swapped from disk into active memory during the interval in kilobyte per second.</p> <p>Key: mem swpinRate_average</p>
Mem Swap Out Rate (KBps)	<p>Rate at which memory is being swapped from active memory to disk during the current interval in kilobytes per second.</p> <p>Key: mem swapoutRate_average</p>
Mem Active Write (KB)	<p>Average active writes in kilobytes.</p> <p>Key: mem activewrite_average</p>
Mem Compressed (KB)	<p>Average memory compression in kilobytes.</p> <p>Key: mem compressed_average</p>
Mem Compression Rate (KBps)	<p>Average compression rate in kilobytes per second.</p> <p>Key: mem compressionRate_average</p>
Mem Decompression Rate (KBps)	<p>Decompression rate in kilobytes per second.</p> <p>Key: mem decompressionRate_average</p>

Metric Name	Description
Mem Total Capacity (KB)	Total capacity in kilobytes. Amount of physical memory configured on the ESXi hosts. Key: mem host_provisioned
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 ISwapInRate_average
Mem Swap In from Host Cache	Amount of memory swapped-in from host cache. Key: mem ISwapIn_average
Mem Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem ISwapOutRate_average
Mem Swap Out to Host Cache	Amount of memory swapped-out to host cache. Key: mem ISwapOut_average
Mem Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem ISwapUsed_average
Mem Low Free Threshold	Threshold of free host physical memory below which ESX begins to reclaim 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
Mem Utilization	Memory utilization level based on descendant Virtual Machines utilization. Includes reservations, limits and overhead to run Virtual Machines Key: mem total_need

Network Metrics for Host Systems

Network metrics provide information about network performance.

Metric Name	Description
Network Driver	This metric displays the type of network driver. Key: net driver
Network Speed	This metric displays the network speed. Key: net speed
Network Management Address	This metric displays the management address of the host network. Key: net management address
Network IP Address	This metric displays the IP address of the host network. Key: net IPaddress
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 Total Throughput (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 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
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 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 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

Metric Name	Description
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
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 Name	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

Metric Name	Description
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
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

Metric Name	Description
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 Name	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 Adapter Metrics for Host Systems

Storage adapter metrics provide information about data storage use.

Metric Name	Description
Storage Adapter Driver	Displays the driver details of the storage adapter. Key: storage adapter driver
Storage Adapter Port WWN	Displays the world wide network port for the storage adapter. Key: storage adapter portwwn
Storage Adapter Total Usage (KBps)	Total latency. Key: storageAdapter usage
Storage Adapter Total IOPS	Average number of commands issued per second by the storage adapter during the collection interval. Key: storageAdapter commandsAveraged_average
Storage Adapter Read IOPS	Average number of read commands issued per second by the storage adapter during the collection interval. Key: storageAdapter numberReadAveraged_average

Metric Name	Description
Storage Adapter Write IOPS	<p>Average number of write commands issued per second by the storage adapter during the collection interval.</p> <p>Key: storageAdapter numberWriteAveraged_average</p>
Storage Adapter Read Throughput (KBps)	<p>Rate of reading data by the storage adapter.</p> <p>Key: storageAdapter read_average</p>
Storage Adapter Read Latency (ms)	<p>This metric shows the average amount of time for a read operation by the storage adapter.</p> <p>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.</p> <p>Total latency is the sum of kernel latency and device latency.</p> <p>Key: storageAdapter totalReadLatency_average</p>
Storage Adapter Write Latency (ms)	<p>This metric shows the average amount of time for a write operation by the storage adapter.</p> <p>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.</p> <p>Total latency is the sum of kernel latency and device latency.</p> <p>Key: storageAdapter totalWriteLatency_average</p>
Storage Adapter Write Throughput (KBps)	<p>Rate of writing data by the storage adapter.</p> <p>Key: storageAdapter write_average</p>
Storage Adapter Demand	<p>Demand.</p> <p>Key: storageAdapter demand</p>
Storage Adapter Highest Latency	<p>Highest Latency.</p> <p>Key: storageAdapter maxTotalLatency_latest</p>
Storage Adapter Outstanding Requests	<p>Outstanding Requests.</p> <p>Key: storageAdapter outstandingIOs_average</p>
Storage Adapter Queue Depth	<p>Queue Depth.</p> <p>Key: storageAdapter queueDepth_average</p>
Storage Adapter Queue Latency (ms)	<p>The average time spent in the ESX Server VM Kernel queue per command.</p> <p>Key: storageAdapter queueLatency_average</p>
Storage Adapter Queued	<p>Queued.</p> <p>Key: storageAdapter queued_average</p>

Metric Name	Description
Storage Adapter Peak Adapter Read Latency	The highest read latency among the storage adapters. A high number indicates that one or more storage adapters are experiencing poor performance. Key: storageAdapter peak_adapter_readLatency
Storage Adapter Peak Adapter Write Latency	The highest write latency among the storage adapters. A high number indicates that one or more storage adapters are experiencing poor performance. Key: storageAdapter peak_adapter_writeLatency

Storage Metrics for Host Systems

Storage metrics provide information about storage use.

Metric Name	Description
Storage Total IOPS	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 Throughput (KBps)	Read throughput rate in kilobytes. Key: storage read_average
Storage Read IOPS	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 Throughput (KBps)	Write throughput rate in kilobytes per second. Key: storage write_average
Storage Write IOPS	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 Name	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 Name	Description
Power Energy (Joule)	Total energy used since last stats reset. 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 Name	Description
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 Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace total_usage
Diskspace Total disk spacey	Total disk space on all datastores visible to this object. Key: diskspace total_capacity

Metric Name	Description
Diskspace Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace total_provisioned .
Diskspace Utilization (GB)	Storage space utilized on connected vSphere datastores. Key: diskspace total_usage
Diskspace Workload (%)	Total storage space available on connected vSphere datastores. Key: diskspace total_capacity

Summary Metrics for Host Systems

Summary metrics provide information about overall host system performance.

Metric Name	Description
Summary Number of Running VMs	<p>This metric shows the number of VMs running on the host during the last metric collection time.</p> <p>Large spikes of running VMs might be a reason for CPU or memory spikes as more resources are used in the host.</p> <p>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 become unavailable (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.</p> <p>Key: summary number_running_vms</p>
Summary Maximum Number of VMs	<p>Maximum number of virtual machines</p> <p>Key: summary max_number_vms</p>
Summary Number of vMotions	<p>This metric shows the number of vMotions that occurred in the host in the last X minutes.</p> <p>The number of vMotions is a good indicator of stability. In a healthy environment, this number should be stable and relatively low.</p> <p>Look for correlation between vMotions and spikes in other metrics such as CPU/Memory contention.</p> <p>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.</p> <p>Key: summary number_vmotion</p>

Metric Name	Description
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
Summary Total Number of VMs	Total number of virtual machines. Note This is the total number of VMs excluding VM templates. Key: summary total_number_vms
Summary Number of VM Templates	Number of VM Templates Key: summary number_vm_templates
Summary Consider for Balance	Summary Consider for Balance = 1 when the host is Powered On, Connected, not in Maintenance Mode, and not a Failover Host, otherwise it = -1

HBR Metrics for Host Systems

Host-based replication (HBR) metrics provide information about vSphere replication.

Metric Name	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

Cost Metrics for Host Systems

Cost metrics provide information about the cost.

Metric Name	Description
Monthly Maintenance Total Cost	Monthly total cost for maintenance. Key: cost maintenanceTotalCost
Monthly Host OS License Total Cost	Monthly total cost for the host operating system license. Key: cost hostOsTotalCost
Monthly Network Total Cost	Monthly total cost for network including cost of NIC cards associated with host. Key: cost networkTotalCost
Monthly Server Hardware Total Cost	Monthly total cost for server hardware, based on amortized monthly value. Key: cost hardwareTotalCost

Metric Name	Description
Monthly Facilities Total Cost	Monthly total cost of facilities including real estate, power, and cooling. Key: cost facilitiesTotalCost
Monthly Server Labor Total Cost	Monthly total cost for the server operating system labor. Key: cost hostLaborTotalCost
Monthly Server Fully Loaded Cost	Monthly cost for a fully loaded server incorporating all cost driver values attributed to the server. Key: cost totalLoadedCost
MTD Server Total Cost	Month to date cost for a fully loaded server incorporating all cost driver values attributed to the server. Key: totalMTDCost
Server Accumulated Depreciation	Month to date accumulated cost for a deprecated server. Key: Cost Server Accumulated Depreciation
Aggregated Daily Total Cost	Daily aggregate daily total cost of the deleted VM present in the host system. Key: Cost aggregatedDailyTotalCost
Aggregated Deleted VM Daily Total Cost	Daily aggregate cost of the deleted VM present in the host system. Key: Cost aggregatedDeletedVmDailyTotalCost

Disabled Instanced Metrics

The instance metrics created for the following metrics are disabled in this version of vRealize Operations Manager . This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

Metric Name
Datastore Outstanding IO requests (OIOs)
Datastore Read IOPS
Datastore Read Latency (ms)
Datastore Read Throughput (KBps)
Datastore Total Latency (ms)
Datastore Total Throughput (KBps)
Datastore unmapIOs_summation
Datastore unmapsize_summation
Datastore Write IOPS
Datastore Write Latency (ms)
Datastore Write Throughput (KBps)

Metric Name
Disk Physical Device Latency (ms)
Disk Queue Latency (ms)
Disk Read IOPS
Disk Read Latency (ms)
Disk Read Throughput (KBps)
Disk Write IOPS
Disk Write Latency (ms)
Disk Write Throughput (KBps)
Net Data Receive Rate (KBps)
Net Data Transmit Rate (KBps)
Net Error Packets Transmitted
Net Packets Dropped (%)
Net Packets Transmitted per second
Net Received Packets Dropped
Net Transmitted Packets Dropped
Net Usage Rate (%)
Storage Adapter Read IOPS
Storage Adapter Read Latency (ms)
Storage Adapter Read Throughput (KBps)
Storage Adapter Write IOPS
Storage Adapter Write Latency (ms)
Storage Adapter Write Throughput (KBps)

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Key
CPU Idle (msec)	cpulidle_summation
CPU Used (msec)	cpulused_summation
Datastore I/O Average Observed Virtual Machine Disk I/O Workload	datastore vmPopulationAvgWorkload
Datastore I/O Max Observed Number of Outstanding IO Operations	datastore maxObserved_OIO
Datastore I/O Max Observed Read Rate (kbps)	datastore maxObserved_Read
Datastore I/O Max Observed Reads per second	datastore maxObserved_NumberRead
Datastore I/O Max Observed Write Rate (kbps)	datastore maxObserved_Write
Datastore I/O Max Observed Writes per second	datastore maxObserved_NumberWrite
Datastore I/O Maximum Observed VM Disk I/O Workload	datastore vmPopulationMaxWorkload
Network I/O bytesRx (kbps)	net bytesRx_average
Network I/O bytesTx (kbps)	net bytesTx_average
Network I/O Demand (%)	net demand
Network I/O Error Packets Received	net errorsRx_summation
Network I/O Max Observed Received Throughput (kbps)	net maxObserved_Rx_KBps
Network I/O Max Observed Throughput (kbps)	net maxObserved_KBps
Network I/O Max Observed Transmitted Throughput (kbps)	net maxObserved_Tx_KBps
Network I/O Packets Received per second	net packetsRxPerSec
Network I/O Packets Dropped	net dropped
Summary Workload Indicator	summary workload_indicator
vFlash Module Latest Number of Active Vm Disks	vflashModule numActiveVMDKs_latest
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 Name	Key
Net Packets Dropped (%)	<p>This metric shows the percentage of received and transmitted packets dropped during the collection interval.</p> <p>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.</p> <p>Key: net droppedPct</p>
Diskspace Not Shared (GB)	<p>Unshared disk space in gigabytes.</p> <p>Key: diskspace notshared</p>

Cluster Compute Resource Metrics

vRealize Operations Manager collects configuration, 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 Analytics Generated Metrics](#)
- [Badge Metrics](#)

Configuration Metrics for Cluster Compute Resources

Configuration metrics provide information about configuration settings.

Metric Name	Description
Configuration DAS Configuration Admission Control Enabled	<p>DAS configuration admission control enabled.</p> <p>Key: configuration dasconfig AdministrationControlEnabled</p>
Configuration DAS Configuration Active Admission Control Policy	<p>DAS configuration active admission control policy.</p> <p>Key: configuration dasconfig activeAdministrationControlPolicy</p>
Configuration DRS Configuration Affinity Rules	<p>Affinity rules for DRS configuration.</p> <p>Key: configuration DRSconfiguration affinity rules</p>
Configuration DRS Configuration Tolerance Imbalance Threshold	<p>Displays the tolerance imbalance threshold for DRS configuration.</p> <p>Key: configuration DRSconfiguration ToleranceimbalanceThreshold</p>
Configuration DRS Configuration Default DRS behavior	<p>Displays the default DRS configuration behavior.</p> <p>Key: configuration DRSconfiguration DefaultDRSbehaviour</p>
Configuration DRS Configuration Idle Consumed Memory	<p>Displays the idle memory consumed by DRS configuration.</p> <p>Key: configuration DRSconfiguration IdleConsumedMemory</p>
Configuration DRS Configuration DRS vMotion Rate	<p>Displays the vMotion rate for the DRS configuration.</p> <p>Key: configuration DRSconfiguration DRSvMotion Rate</p>
Configuration DPM Configuration Default DPM behavior	<p>Displays the default behavior for the DPM configuration.</p> <p>Key: configuration DPMconfiguration DefaultDPMbehaviour</p>

Metric Name	Description
Configuration DPM Configuration DPM Enabled	Displays whether the DPM Configuration is enabled or not. Key: configuration DPMConfiguration DPMEnabled
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
Configuration CPU Failover Resources Percent	Percent CPU failover resources for DAS configuration admission control policy. Key: configuration dasconfig admissionControlPolicy cpuFailoverResourcesPercent
Configuration Memory Failover Resources Percent	Percent memory failover resources for DAS configuration admission control policy. Key: configuration dasconfig admissionControlPolicy memoryFailoverResourcesPercent

Disk Space Metrics for Cluster Compute Resources

Disk space metrics provide information about disk space use.

Metric Name	Description
DiskSpace Snapshot Space	Displays the disk space used by the snapshot. Key: DiskSpace snapshot space
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 Utilization (GB)	Storage space used on the connected vSphere Datastores. Key: diskSpace total_usage
DiskSpace Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: diskSpace total_capacity

CPU Usage Metrics for Cluster Compute Resources

CPU usage metrics provide information about CPU use.

Metric Name	Description
CPU Allocation Usable Capacity after HA and Buffer (vCPUs)	<p>This metric shows the total capacity taking into consideration the over-commit ratio and after subtracting the CPU resources needed for HA and reserved buffer.</p> <p>Key: cpu alloc usableCapacity</p>
CPU Capacity Usage	<p>This metric shows the percentage of the capacity used.</p> <p>Key: cpu capacity_usagepct_average</p>
CPU CPU Contention (%)	<p>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.</p> <p>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:</p> <ul style="list-style-type: none"> ■ CPU Ready ■ CPU Co-stop ■ Power management ■ Hyper threading <p>This metric is more accurate than CPU Ready since it takes into account CPU Co-stop and Hyper threading.</p> <p>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%.</p> <p>Since this value is averaged across all hosts in the cluster, you might find that some hosts have a higher CPU contention while others are lower. To ensure that vSphere spreads out the running workloads across hosts, consider enabling a fully automated DRS in the cluster.</p> <p>Key: cpu capacity_contentionPct</p>
CPU Demand Usable Capacity after HA and Buffer (MHz)	<p>This metric shows the total capacity after subtracting the CPU resources needed for HA and reserved buffer.</p> <p>Key: cpu demand usableCapacity</p>
CPU Demand (%)	<p>This metric is an indicator of the overall demand for CPU resources by the workloads in the cluster.</p> <p>It shows the percentage of CPU resources that all the virtual machines might use if there were no CPU contention or CPU limits set. It represents the average active CPU load in the past five minutes.</p> <p>Key: cpu demandPct</p>

Metric Name	Description
CPU Demand (MHz)	CPU utilization level based on descendant virtual machines utilization. This Includes reservations, limits, and overhead to run the virtual machines. Key: cpuldemandmhz
CPU Number of CPU Sockets	Number of CPU sockets. Key: cpulnumpackages
CPU Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpulcapacity_contention
CPU Host Provisioned Capacity	Provisioned CPU capacity in megahertz. Key: cpulcapacity_provisioned
CPU Provisioned vCPUs	Number of provisioned CPU cores. Key: cpulcorecount_provisioned
CPU Usage (MHz)	Average CPU use in megahertz. Key: cpulusagemhz_average
CPU Demand	CPU Demand. Key: cpuldemand_average
CPU Overhead	Amount of CPU overhead. Key: cpuloverhead_average
CPU Demand without overhead	Value of demand excluding any overhead. Key: cpuldemand_without_overhead
CPU Provisioned Capacity	Provisioned Capacity (MHz). Key: cpulvm_capacity_provisioned
CPU Number of hosts stressed	Number of hosts stressed. Key: cpulnum_hosts_stressed
CPU Stress Balance Factor	Stress Balance Factor. Key: cpulstress_balance_factor
CPU Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: cpulmin_host_capacity_remaining
CPU Workload Balance Factor	Workload Balance Factor. Key: cpulworkload_balance_factor
CPU Highest Provider Workload	Highest Provider Workload. Key: cpulmax_host_workload
CPU Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: cpulhost_workload_disparity
CPU Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: cpulhost_stress_disparity

Metric Name	Description
CPU Total Capacity (MHz)	Total CPU resources configured on the descendant ESXi hosts. Key: cpu capacity_provisioned
CPU Usable Capacity (MHz)	The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services. Key: cpu haTotalCapacity_average

Disk Metrics for Cluster Compute Resources

Disk metrics provide information about disk use.

Metric Name	Description
Disk Total IOPS	Average number of commands issued per second during the collection interval. Key: disk commandsAveraged_average
Disk Total 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 Read Latency (ms)	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 Write 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 totalWriteLatency_averag
Disk Read IOPS	Average number of read commands issued per second during the collection interval. Key: disk numberReadAveraged_averag
Disk Total Throughput (KBps)	Average of the sum of the data read and written for all the disk instances of the host or virtual machine. Key: disk usage_average
Disk Write IOPS	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

Metric Name	Description
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 Name	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 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

Metric Name	Description
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
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 swpin_average
Mem Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem swpinRate_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

Metric Name	Description
Mem Usage (%)	<p>This metric shows the portion of the total memory in all hosts in the cluster that is being used.</p> <p>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.</p> $\frac{\sum \text{memory consumed on all hosts}}{\sum \text{physical memory on all hosts}} \times 100\%$
Mem Usage (KB)	<p>Memory currently in use as a percentage of total available memory.</p> <p>Key: mem usage_average</p>
Mem VM kernel Usage (KB)	<p>Amount of memory that the VM kernel uses.</p> <p>Key: mem sysUsage_average</p>
Mem Zero (KB)	<p>Amount of memory that is all 0.</p> <p>Key: mem zero_average</p>
Mem Number of Hosts Stressed	<p>Number of hosts stressed.</p> <p>Key: mem num_hosts_stressed</p>
Mem Stress Balance Factor	<p>Stress balance factor.</p> <p>Key: mem stress_balance_factor</p>
Mem Lowest Provider Capacity Remaining	<p>Lowest provider capacity remaining.</p> <p>Key: mem min_host_capacity_remaining</p>
Mem Workload Balance Factor	<p>Workload balance factor.</p> <p>Key: mem workload_balance_factor</p>
Mem Highest Provider Workload	<p>Highest provider workload.</p> <p>Key: mem max_host_workload</p>
Mem Host workload Max-Min Disparity	<p>Difference of Max and Min host workload in the container.</p> <p>Key: mem host_workload_disparity</p>
Mem Host stress Max-Min Disparity	<p>Difference of Max and Min host stress in the container.</p> <p>Key: mem host_stress_disparity</p>
Mem Utilization (KB)	<p>Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines.</p> <p>Key: mem total_need</p>
Mem Total Capacity (KB)	<p>Total physical memory configured on descendant ESXi hosts.</p> <p>Key: mem host_provisioned</p>
Mem Usable Capacity (KB)	<p>The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services.</p> <p>Key: mem haTotalCapacity_average</p>

Network Metrics for Cluster Compute Resources

Network metrics provide information about network performance.

Metric Name	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
Net Total Throughput (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net usage_average

Datastore Metrics for Cluster Compute Resources

Datastore metrics provide information about Datastore use.

Metric Name	Description
Datastore TotalThroughput	Displays the total throughput for the datastore. Key: datastore thoroughput
Datastore Outstanding IO requests	OIO for datastore. Key: datastore demand_oio
Datastore Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Datastore Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Datastore Read Throughput (KBps)	Amount of data read in the performance interval. Key: datastore read_average
Datastore Write Throughput (KBps)	Amount of data written to disk in the performance interval. Key: datastore write_average

Metric Name	Description
Datastore Read Latency	Average amount of time taken for a read operation from the datastore. Key: datastore ReadLatency
Datastore Write Latency	Average amount of time taken for a write operation from the datastore. Key: datastore WriteLatency
Datastore Max VM Disk Latency	Maximum amount of time taken to read or write data from a virtual machine. Key: datastore MaxVMDiskLatency
Datastore Outstanding IO Requests (OIOs)	This metric displays the outstanding datastore IO requests. Key: datastore OutstandingIORequests
Datastore Host SCSI Disk Partition	This metric displays the datastore host scsi partition. Key: datastore HostSCSIDiskPartition
Devices Command Aborted	The metric lists the stopped device commands. Key: devices CommandAborted

Cluster Services Metrics for Cluster Compute Resources

Cluster Services metrics provide information about cluster services.

Metric Name	Description
Cluster Services Total Imbalance	Total imbalance in cluster services Key: clusterServices total_imbalance
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
Cluster Services DRS Initiated vMotion Count	clusterServices number_drs_vmotion

Power Metrics for Cluster Compute Resources

Power metrics provide information about power use.

Metric Name	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 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 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 might serve its VMs. A vMotion can impact VM performance during the stun time. Key: summary number_vmotion
Summary Number of Hosts	Total number of hosts. Key: summary total_number_hosts
Summary Total Number of VMs	Total number of virtual machines. Note This shows the total number of VMs excluding VM templates under the datastore. Key: summary total_number_vms
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 Cluster Availability (%)	Percentage of hosts powered-on in the cluster. Key: summary cluster_availability
Summary Datastore	Displays the status of the datastore. Key: summary datastore
Summary Type	Displays the datastore type. Key: summary type
Summary Is Local	Displays whether the datastore is local or not. Key: summary islocal
Summary Number of VM Templates	Number of VM templates. Key: summary number_vm_templates
Summary Number of Pods	Number of pods. Note This is published if the cluster is Workload Management enabled or there are pods under the cluster. Key: summary total_number_pods

Metric Name	Description
Summary Number of Namespaces	<p>Number of namespaces.</p> <p>Note This is published if the cluster is Workload Management enabled or there are namespaces under the cluster.</p> <p>Key: summary numberNamespaces</p>
Summary Number Kubernetes Clusters	<p>Number of Kubernetes clusters.</p> <p>Note This is published if the cluster is Workload Management enabled or there are Kubernetes clusters under the cluster.</p> <p>Key: summary numberKubernetesClusters</p>
Summary Number of Developer Managed VMs	<p>Number of developer managed VMs.</p> <p>Note This is published if the cluster is Workload Management enabled or there are developer managed VMs under the cluster.</p> <p>Key: summary numberDeveloperManagedVMs</p>
Namespaces Config Status	<p>Workload Management configuration status.</p> <p>Note This is published if the cluster is Workload Management enabled.</p> <p>Key: namespaces configStatus</p>
Namespaces Kubernetes Status	<p>Kubernetes status.</p> <p>Note This is published if the cluster is Workload Management enabled.</p> <p>Key: namespaces kubernetesStatus</p>

Reclaimable Metrics for Cluster Compute Resources

Reclaimable metrics provide information about reclaimable resources.

Metric Name	Description
Idle VMs CPU (vCPUs)	<p>Number of reclaimable vCPUs of Idle VMs within the cluster.</p> <p>Key: reclaimable idle_vms cpu</p>
Idle VMs Disk Space (GB)	<p>Reclaimable disk space of Idle VMs within the cluster.</p> <p>Key: reclaimable idle_vms diskspace</p>
Idle VMs Memory (KB)	<p>Reclaimable memory of Idle VMs within the cluster.</p> <p>Key: reclaimable idle_vms mem</p>
Idle VMs Potential Savings	<p>Potential saving after reclamation of resources of Idle VMs within the cluster.</p> <p>Key: reclaimable idle_vms cost</p>
Powered Off VMs Disk Space (GB)	<p>Reclaimable disk space of Powered Off VMs within the cluster.</p> <p>Key: reclaimable poweredOff_vms diskspace</p>

Metric Name	Description
Powered Off VMs Potential Savings	Potential saving after reclamation of resources of Powered Off VMs within the cluster. Key: reclaimable poweredOff_vms cost
VM Snapshots Disk Space (GB)	Reclaimable disk space of VM Snapshots within the cluster. Key: reclaimable vm_snapshots diskspace
VM Snapshots Potential Savings	Potential saving after reclamation of VM Snapshots within the cluster. Key: reclaimable vm_snapshots cost

Cost Metrics for Cluster Compute Resources

Cost metrics provide information about the cost.

Metric Name	Description
Cluster CPU Base Rate	Base rate for Cluster CPU calculated by dividing the monthly total cluster CPU cost by cluster CPU utilization % and CPU cluster capacity (GHz). Key: cost cpuBaseRate
Cluster CPU Utilization (%)	Expected CPU utilization that is set by the user in cluster cost page. Key: cost cpuExpectedUtilizationPct
Cluster Memory Base Rate	Cluster memory base rate calculated by dividing the monthly total cluster memory cost by cluster memory utilization % and memory cluster capacity (GB). Key: cost memoryBaseRate
Cluster Memory Utilization (%)	Expected memory utilization that is set by the user in cluster cost page. Key: cost memoryExpectedUtilizationPct
Monthly Cluster Allocated Cost	Monthly cluster allocated cost calculated by subtracting the monthly cluster unallocated cost from the monthly cluster total cost. Key: cost allocatedCost
Monthly Cluster Total Cost	Fully loaded compute cost of all hosts underneath the cluster. Key: cost totalCost
Monthly Cluster Unallocated Cost	Monthly cluster unallocated cost calculated by subtracting the monthly cluster allocated cost from the monthly cluster total cost. Key: cost unAllocatedCost
Monthly Total Cluster CPU Cost	Cost attributed to the cluster CPU from monthly cluster total cost. Key: cost totalCpuCost

Metric Name	Description
Monthly Total Cluster Memory Cost	Cost attributed to the cluster memory from monthly cluster total cost. Key: cost totalMemoryCost
MTD Cluster CPU Utilization (GHz)	Month to date CPU utilization of the cluster. Key: cost cpuActualUtilizationGHz
MTD Cluster Memory Utilization (GB)	Month to date memory utilization of the cluster. Key: cost memoryActualUtilizationGB
Monthly Cluster Allocated Cost (Currency)	The monthly allocated cost of all VMs in a cluster. cost clusterAllocatedCost
Cost Allocation Monthly Cluster Unallocated Cost (Currency)	The monthly unallocated is calculated by subtracting the monthly allocated cost from the cluster's cost. cost clusterUnAllocatedCost
Aggregated Daily Total Cost	Daily aggregate daily total cost of the deleted VM present in the host system. Key: Cost aggregatedDailyTotalCost
Aggregated Deleted VM Daily Total Cost	Daily aggregate cost of the deleted VM present in the host system. Key: Cost aggregatedDeletedVmDailyTotalCost

Profiles Metrics for Cluster Compute Resources

Profiles metrics provide information about the profile specific capacity.

Metric Name	Description
Profiles Capacity Remaining Profile (Average)	The capacity remaining in terms of fitting the average consumer. Key: Profiles capacityRemainingProfile_<profile uuid>
Profiles Capacity Remaining Profile (<custom profile name>)	Published for custom profiles enabled from policy on Cluster Compute Resource. Key: Profiles capacityRemainingProfile_<profile uuid>

Capacity Allocation Metrics for Cluster Compute Resources

Capacity allocation metrics provide information about the allotment of capacity, see [Capacity Analytics Generated Metrics](#).

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see *Metrics and Properties Details*.

Metric Name	Key
CPU Capacity Available to VMs (mhz)	cpu totalCapacity_average
CPU IO Wait (msec)	cpu iowait
CPU Reserved Capacity (mhz)	cpu reservedCapacity_average
CPU Total Wait (msec)	cpu wait
Datastore I/O Max Observed Number of Outstanding IO Operations	datastore maxObserved_OIO
Datastore I/O Max Observed Read Rate (kbps)	datastore maxObserved_Read
Datastore I/O Max Observed Reads per second	datastore maxObserved_NumberRead
Datastore I/O Max Observed Write Rate (kbps)	datastore maxObserved_Write
Datastore I/O Max Observed Writes per second	datastore maxObserved_NumberWrite
Storage Total Usage (kbps)	storage usage_average
Summary Average Provisioned Capacity per Running VM (mhz)	summary avg_vm_cpu
Summary Average Provisioned Memory per Running VM (kb)	summary avg_vm_mem
Summary Average Provisioned Memory per Running VM (kb)	summary avg_vm_mem
Summary Maximum Number of VMs	summary max_number_vms
Summary Workload Indicator	summary workload_indicator
Network I/O Max Observed Received Throughput (KBps)	net maxObserved_Rx_KBps
Network I/O Max Observed Throughput (KBps)	net maxObserved_KBps
Network I/O Max Observed Transmitted Throughput (KBps)	net maxObserved_Tx_KBps
Diskspace Not Shared (GB)	Space used by VMs that is not shared. Key: diskspace notshared

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 Analytics Generated Metrics](#)
- [Badge Metrics](#)

Configuration Metrics for Resource Pools

Configuration metrics provide information about memory and CPU allocation configuration.

Metric Name	Description
Memory Allocation Reservation	Memory Allocation Reservation. Key: config mem_alloc_reservation

CPU Usage Metrics for Resource Pools

CPU usage metrics provide information about CPU use.

Metric Name	Description
Capacity Demand Entitlement (%)	CPU Capacity Demand Entitlement Percentage. Key: cpu capacity_demandEntitlementPct
Capacity entitlement (MHz)	CPU Capacity Entitlement. Key: cpu capacity_entitlement
CPU Contention (%)	CPU capacity contention. Key: cpu capacity_contentionPct
Demand (MHz)	CPU demand in megahertz. Key: cpu demandmhz
Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpu capacity_contention
Usage	Average CPU use in megahertz. Key: cpu usagemhz_average
Effective limit	CPU effective limit. Key: cpu effective_limit
Reservation Used	CPU reservation used. Key: cpu reservation_used
Estimated entitlement	CPU estimated entitlement. Key: cpu estimated_entitlement
Dynamic entitlement	CPU dynamic entitlement. Key: cpu dynamic_entitlement
Demand without overhead	Value of demand excluding any overhead. Key: cpu demand_without_overhead

Memory Metrics for Resource Pools

Memory metrics provide information about memory use and allocation.

Metric Name	Description
Balloon	Amount of memory currently used by the virtual machine memory control. Key: mem vmmemctl_average
Compression Rate	Compression rate in kilobytes per second. Key: mem compressionRate_average
Consumed	Amount of host memory consumed by the virtual machine for guest memory. Key: mem consumed_average
Contention	Machine contention. Key: mem host_contentionPct
Guest usage	Guest memory entitlement. Key: mem guest_usage
Guest demand	Guest memory entitlement. Key: mem guest_demand
Contention (KB)	Machine contention in kilobytes. Key: mem host_contention
Decompression Rate	Decompression rate in kilobytes per second. Key: mem decompressionRate_average
Granted	Average of memory available for use. Key: mem granted_average
Guest Active	Amount of memory that is actively used. Key: mem active_average
VM Overhead	Memory overhead reported by host. Key: mem overhead_average
Shared	Amount of shared memory. Key: mem shared_average
Reservation Used	Memory Reservation Used. Key: mem reservation_used
Dynamic Entitlement	Memory Dynamic Entitlement. Key: mem dynamic_entitlement
Effective Limit	Memory Effective Limit. Key: mem effective_limit
Swap In Rate	Rate at which memory is swapped from disk into active memory during the interval. Key: mem swpinRate_average
Swap Out Rate	Rate at which memory is being swapped from active memory to disk during the current interval. Key: mem swapoutRate_average
Swapped	Amount of unreserved memory. Key: mem swapped_average

Metric Name	Description
Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem usage_average
Zero	Amount of memory that is all zero. Key: mem zero_average
Zipped (KB)	Latest zipped memory in kilobytes. Key: mem zipped_latest
Swap In (KB)	Amount of memory swapped in kilobytes. Key: mem swpin_average
Swap Out (KB)	Amount of memory swapped out in kilobytes. Key: mem swpout_average
Swap Used	Amount of memory used for swap space in kilobytes. Key: mem swpused_average
Total Capacity	Total capacity. Key: mem guest_provisioned

Summary Metrics for Resource Pools

Summary metrics provide information about overall performance.

Metric Name	Description
Number of Running VMs	Number of running virtual machines. Key: summary number_running_vms
Total Number of VMs	Total number of virtual machines. Note This shows the total number of VMs excluding VM templates. Key: summary total_number_vms
IO Wait (ms)	IO wait time in milliseconds. Key: summary iowait
Number of VM Templates	Number of VM Templates. Key: summary number_vm_templates

Data Center Metrics

vRealize Operations Manager collects CPU usage, disk, memory, network, storage, disk space, and summary metrics for data center objects.

Data center metrics include capacity and badge metrics. See definitions in:

- [Capacity Analytics Generated Metrics](#)
- [Badge Metrics](#)

CPU Usage Metrics for Data Centers

CPU usage metrics provide information about CPU use.

Metric Name	Description
Capacity Usage (%)	Percent capacity used. Key: <code>cpu capacity_usagepct_average</code>
CPU Contention (%)	CPU capacity contention. Key: <code>cpu capacity_contentionPct</code>
Demand (%)	CPU demand percentage. Key: <code>cpu demandPct</code>
Demand	Demand in megahertz. Key: <code>cpu demandmhz</code>
Demand (MHz)	CPU utilization level based on descendant virtual machines utilization. This includes reservations, limits, and overhead to run the virtual machines. Key: <code>cpu demandmhz</code>
Overhead (KB)	Amount of CPU overhead. Key: <code>cpu overhead_average</code>
Demand without overhead	Value of demand excluding any overhead. Key: <code>cpu demand_without_overhead</code>
Total Wait	CPU time spent on idle state. Key: <code>cpu wait</code>
Number of CPU Sockets	Number of CPU sockets. Key: <code>cpu numpackages</code>
Overall CPU Contention (ms)	Overall CPU contention in milliseconds. Key: <code>cpu capacity_contention</code>
Host Provisioned Capacity (MHz)	Host provisioned capacity in megahertz. Key: <code>cpu capacity_provisioned</code>
Provisioned vCPU(s)	Provisioned vCPU(s). Key: <code>cpu corecount_provisioned</code>
Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool. Key: <code>cpu reservedCapacity_average</code>
Usage	Average CPU usage in megahertz. Key: <code>cpu usagemhz_average</code>
IO Wait	IO wait time in milliseconds. Key: <code>cpu iowait</code>
Provisioned Capacity	Provisioned Capacity. Key: <code>cpu vm_capacity_provisioned</code>
Stress Balance Factor	Stress Balance Factor. Key: <code>cpu stress_balance_factor</code>

Metric Name	Description
Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: cpulmin_host_capacity_remaining
Workload Balance Factor	Workload Balance Factor. Key: cpu workload_balance_factor
Highest Provider Workload	Highest Provider Workload. Key: cpu max_host_workload
Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: cpu host_workload_disparity
Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: cpu host_stress_disparity
Total Capacity (MHz)	Total CPU resources configured on the descendant ESXi hosts. Key: cpu capacity_provisioned
Usable Capacity (MHz)	The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services. Key: cpu haTotalCapacity_average

Disk Metrics for Data Centers

Disk metrics provide information about disk use.

Metric Name	Description
Total IOPS	Average number of commands issued per second during the collection interval. Key: disk commandsAveraged_average
Total 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 Latency and Physical Device Latency metrics. Key: disk totalLatency_average
Total Throughput (KBps)	Average of the sum of the data read and written for all the disk instances of the host or virtual machine. Key: disk usage_average
Total queued outstanding operations	Sum of queued operations and outstanding operations. Key: disk sum_queued_oio
Max observed OIO	Max observed IO for a disk. Key: disk max_observed

Memory Metrics for Data Centers

Memory metrics provide information about memory use and allocation.

Metric Name	Description
Contention (%)	Machine Contention Percentage. Key: mem host_contentionPct
Machine Demand (KB)	Memory machine demand in kilobytes. Key: mem host_demand
ESX System Usage	Memory usage by the VM kernel and ESX user-level services. Key: mem host_systemUsage
Provisioned Memory (KB)	Provisioned host memory in kilobytes. Key: mem host_provisioned
Reserved Capacity (KB)	Reserved memory capacity in kilobytes. Key: mem reservedCapacity_average
Usable Memory (KB)	Usable host memory in kilobytes. Key: mem host_usable
Host Usage	Host memory use in kilobytes. Key: mem host_usage
Usage/Usable (%)	Percent host memory used. Key: mem host_usagePct
VM Overhead	Memory overhead reported by host. Key: mem overhead_average
Stress Balance Factor	Stress Balance Factor. Key: mem stress_balance_factor
Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: mem min_host_capacity_remaining
Workload Balance Factor	Workload Balance Factor. Key: mem workload_balance_factor
Highest Provider Workload	Highest Provider Workload. Key: mem max_host_workload
Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: mem host_workload_disparity
Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: mem host_stress_disparity
Utilization (KB)	Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines. Key: mem total_need

Metric Name	Description
Total Capacity (KB)	Total physical memory configured on descendant ESXi hosts. Key: mem host_provisioned
Usable Capacity (KB)	The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services. Key: mem haTotalCapacity_average

Network Metrics for Data Centers

Network metrics provide information about network performance.

Metric Name	Description
Packets Dropped	Percentage of packets dropped. Key: net droppedPct
Max Observed Throughput	Max observed rate of network throughput. Key: net maxObservedKBps
Data Transmit Rate	Average amount of data transmitted per second. Key: net transmitted_average
Data Receive Rate	Average amount of data received per second. Key: net received_average
Total Throughput (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net usage_average

Storage Metrics for Data Centers

Storage metrics provide information about storage use.

Metric Name	Description
Total Usage	Total throughput rate. Key: storage usage_average

Datastore Metrics for Data Centers

Datastore metrics provide information about Datastore use.

Metric Name	Description
Outstanding IO requests	OIO for datastore. Key: datastore demand_oio
Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average

Metric Name	Description
Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Read Throughput (KBps)	Amount of data read in the performance interval. Key: datastore read_average
Write Throughput (KBps)	Amount of data written to disk in the performance interval. Key: datastore write_average

Disk Space Metrics for Data Centers

Disk space metrics provide information about disk use.

Metric Name	Description
Virtual machine used	Used virtual machine disk space in gigabytes. Key: diskspace used
Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace total_usage
Total disk space	Total disk space on all datastores visible to this object. Key: diskspace total_capacity
Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace total_provisioned
Shared Used (GB)	Shared disk space in gigabytes. Key: diskspace shared
Snapshot Space (GB)	Snapshot disk space in gigabytes. Key: diskspace snapshot
Virtual Disk Used (GB)	Used virtual disk space in gigabytes. Key: diskspace diskused
Number of Virtual Disks	Number of Virtual Disks. Key: diskspace numvmdisk
Utilization (GB)	Storage space used on the connected vSphere Datastores. Key: diskspace total_usage
Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: diskspace total_capacity

Summary Metrics for Data Centers

Summary metrics provide information about overall performance.

Metric Name	Description
Number of Running Hosts	Number of hosts that are ON. Key: summary number_running_hosts
Number of Running VMs	Number of running virtual machines. Key: summary number_running_vms
Maximum Number of VMs	Maximum number of virtual machines. Key: summary max_number_vms
Number of Clusters	Total number of clusters. Key: summary total_number_clusters
Number of Hosts	Total number of hosts. Key: summary total_number_hosts
Number of VMs	Total number of virtual machines. Key: summary total_number_vms
Total Number of Datastores	Total number of datastores. Key: summary total_number_datastores
Number of VCPUs on Powered On VMs	Total number of VCPUs of virtual machines that are powered on. Key: summary number_running_vcpus
Workload Indicator	Workload indicator. Key: summary workload_indicator
Average Running VM Count per Running Host	Average number of running virtual machines per running host. Key: summary avg_vm_density

Reclaimable Metrics for Data Centers

Reclaimable metrics provide information about reclaimable resources.

Metric Name	Description
CPU (vCPUs)	Number of reclaimable vCPUs within the data center. Key: reclaimable cpu
Disk Space	Reclaimable disk space within the data center. Key: reclaimable diskspace
Potential Savings	Potential saving after reclamation of resources of all reclaimable VMs (Idle VMs, Powered Off VMs, VM snapshots) within the data center. Key: reclaimable cost
Memory (KB)	Reclaimable memory within the data center. Key: reclaimable mem
Virtual Machines	Number of VMs having reclaimable resources (Memory, disk space, vCPU) within the data center. Key: reclaimable vm_count

Metric Name	Description
Idle VMs Potential Savings	Potential saving after reclamation of resources of Idle VMs within the data center. Key: reclaimable idle_vms cost
Powered Off VMs Potential Savings	Potential saving after reclamation of resources of Powered Off VMs within the data center. Key: reclaimable poweredOff_vms cost
VM Snapshots Potential Savings	Potential saving after reclamation of VM snapshots within the data center. Key: reclaimable vm_snapshots cost
Reclaimable Orphaned Disks Potential Savings (Currency)	Displays the potential savings after reclamation of disk space by removing orphaned VMDks from all datastores under datacenter. reclaimable cost
Reclaimable Number of Orphaned Disks	Number of reclaimable orphaned disks is the sum of all orphaned disks on it's datastore. reclaimable orphaned_disk_count

Cost Metrics for Data Centers

Cost metrics provide information about the cost.

Metric Name	Description
Monthly Cluster Aggregated Allocated Cost	Sum of the monthly allocated cost for both cluster and unclustered hosts. Key: cost clusterAllocatedCost
Monthly Cluster Aggregated Cost	The sum of monthly aggregated allocated and unallocated cost for both cluster and unclustered hosts. Key: cost clusterCost
Monthly Cluster Aggregated Unallocated Cost	Sum of the monthly unallocated cost for both cluster and unclustered hosts. Key: cost clusterUnAllocatedCost
Monthly Datacenter Aggregated Total Cost	Monthly aggregated total cost for the data center. Key: cost aggrTotalCost
Monthly Datastore Total Cost	Monthly data store total cost. Key: cost totalCost
Monthly Datastore Aggregated Allocated Cost	Monthly aggregated allocated cost for the datastore. Key: cost aggrDataStoreAllocatedCost
Monthly Datastore Aggregated Unallocated Cost	Monthly aggregated unallocated cost for the datastore. Key: cost aggrDataStoreUnallocatedCost
Monthly VM Aggregated Direct Cost	Month to date aggregated VM direct cost across all the VMs under the data center. Key: cost vmDirectCost

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Key
Datastore I/O Max Observed Number of Outstanding IO Operations (IOPS)	datastore maxObserved_OIO
Datastore I/O Max Observed Read Rate (KBps)	datastore maxObserved_Read
Datastore I/O Max Observed Reads per second (IOPS)	datastore maxObserved_NumberRead
Datastore I/O Max Observed Write Rate (KBps)	datastore maxObserved_Write
Datastore I/O Max Observed Writes per second (IOPS)	datastore maxObserved_NumberWrite
Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net maxObserved_Tx_KBps
Max Observed Received Throughput	Max observed received rate of network throughput. Key: net maxObserved_Rx_KBps
Not Shared (GB)	Unshared disk space in gigabytes. Key: diskspace notshared

Custom Data Center Metrics

vRealize Operations Manager collects CPU usage, memory, summary, network, and datastore metrics for custom data center objects.

Custom data center metrics include capacity and badge metrics. See definitions in:

- [Capacity Analytics Generated Metrics](#)
- [Badge Metrics](#)

CPU Usage Metrics for Custom Data Centers

CPU usage metrics provide information about CPU use.

Metric Name	Description
Host Provisioned Capacity	Host provisioned capacity (MHz). Key: cpu capacity_provisioned
Provisioned vCPU(s)	Provisioned vCPU(s). Key: cpu corecount_provisioned
Demand without overhead	Value of demand excluding any overhead. Key: cpudemand_without_overhead
Number of hosts stressed	Number of hosts stressed. Key: cpu num_hosts_stressed

Metric Name	Description
Stress Balance Factor	Stress balance factor. Key: cpu stress_balance_factor
Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: cpu min_host_capacity_remaining
Workload Balance Factor	Workload balance factor. Key: cpu workload_balance_factor
Highest Provider Workload	Highest provider workload. Key: cpu max_host_workload
Host workload Max-Min Disparity	Host workload max-min disparity. Key: cpu host_workload_disparity
Host stress Max-Min Disparity	Difference of max and min host stress in the container. Key: cpu host_stress_disparity
Demand (MHz)	CPU utilization level based on descendant virtual machines utilization. This includes reservations, limits, and overhead to run the virtual machines. Key: cpudemandmhz
Total Capacity (MHz)	Total CPU resources configured on the descendant ESXi hosts. Key: cpu capacity_provisioned
Usable Capacity (MHz)	The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services. Key: cpu haTotalCapacity_average

Memory Metrics for Custom Data Centers

Memory metrics provide information about memory use.

Metric Name	Description
Usable Memory	Usable memory. Key: mem host_usable
Machine Demand	Memory machine demand in KB. Key: mem host_demand
Number of hosts stressed	Number of hosts stressed. Key: mem num_hosts_stressed
Stress Balance Factor	Stress balance factor. Key: mem stress_balance_factor
Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: mem min_host_capacity_remaining
Workload Balance Factor	Workload balance factor. Key: mem workload_balance_factor

Metric Name	Description
Highest Provider Workload	Highest provider workload. Key: mem max_host_workload
Host workload Max-Min Disparity	Host workload max-min disparity. Key: mem host_workload_disparity
Host stress max-min disparity	Host stress max-min disparity. Key: mem host_stress_disparity
Utilization (KB)	Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines. Key: mem total_need
Total Capacity (KB)	Total physical memory configured on descendant ESXi hosts. Key: mem host_provisioned
Usable Capacity (KB)	The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services. Key: mem haTotalCapacity_average

Summary Metrics for Custom Data Centers

Summary metrics provide information about overall performance.

Metric Name	Description
Number of Running VMs	Number of virtual machines that are ON. Key: summary number_running_vms
Maximum Number of VMs	Maximum number of virtual machines. Key: summary max_number_vms
Status	Status of the data center. Key: summary status

Network Metrics for Custom Data Centers

Network metrics provide information about network performance.

Metric Name	Description
Usage Rate	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net usage_average
Data Transmit Rate	Average amount of data transmitted per second. Key: net transmitted_average
Data REceive Rate	Average amount of data received per second. Key: net received_average

Datastore Metrics for Custom Data Centers

Datastore metrics provide information about datastore use.

Metric Name	Description
Outstanding IO requests	OIO for datastore. Key: datastore demand_oio
Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Write IOPS	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Read Throughput (KBps)	Amount of data read in the performance interval. Key: datastore read_average
Write Throughput (KBps)	Amount of data written to disk in the performance interval. Key: datastore write_average

Reclaimable Metrics for Custom Data Centers

Reclaimable metrics provide information about reclaimable resources.

Metric Name	Description
CPU (vCPUs)	Number of reclaimable vCPUs within the custom data center. Key: reclaimable cpu
Disk Space	Reclaimable disk space within the custom data center. Key: reclaimable diskspace
Potential Savings	Potential saving after reclamation of resources of all reclaimable VMs (Idle VMs, Powered Off VMs, VM snapshots) within the custom data center. Key: reclaimable cost
Memory (KB)	Reclaimable memory within the custom data center. Key: reclaimable mem
Number of Orphaned Disks	Number of reclaimable orphaned disks within the custom data center. reclaimable orphaned_disk_count
Reclaimable Orphaned Disks Potential Savings	Potential savings in cost after reclamation of orphaned disks across the custom data center. Key: reclaimable orphaned_disk cost
Note The orphaned disk reclamation feature might not work as expected when vRealize Operations Manager monitors multiple vCenters which use shared data stores.	

Metric Name	Description
Virtual Machines	Number of VMs having reclaimable resources (Memory, disk space, vCPU) within the custom data center. Key: reclaimable vm_count
Idle VMs Potential Savings	Potential saving after reclamation of resources of Idle VMs within the custom data center. Key: reclaimable idle_vms cost
Powered Off VMs Potential Savings	Potential saving after reclamation of resources of Powered Off VMs within the custom data center. Key: reclaimable poweredOff_vms cost
VM Snapshots Potential Savings	Potential saving after reclamation of VM snapshots within the custom data center. Key: reclaimable vm_snapshots cost
Reclaimable Orphaned Disks Potential Savings (Currency)	Displays the potential savings after reclamation of disk space by removing orphaned VMDks from all datastores under custom datacenters. reclaimable cost
Reclaimable Number of Orphaned Disks	Number of reclaimable orphaned disks is the sum of the numbers of orphaned disks on it's datastore. reclaimable orphaned_disk_count

Disk Space Metrics for Custom Data Centers

Disk space metrics provide information about disk use.

Metric Name	Description
Utilization (GB)	Storage space used on the connected vSphere Datastores. Key: diskspace total_usage
Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: diskspace total_capacity

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Key
Max Observed Throughput	Max observed rate of network throughput. Key: net maxObserved_KBps
Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net maxObserved_Tx_KBps

Metric Name	Key
Max Observed Received Throughput	Max observed received rate of network throughput. Key: net maxObserved_Rx_KBps
Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore maxObserved_NumberRead
Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore maxObserved_Read
Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore maxObserved_NumberWrite
Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore maxObserved_Write
Max Observed Number of Outstanding IO Operations	Max observed number of outstanding IO operations. Key: datastore maxObserved_OIO

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 Analytics Generated Metrics](#)
- [Badge Metrics](#)

Table 8-2. Datastore Metrics for Storage Pods

Metric Name	Description
Read IOPS	Average number of read commands issued per second during the collection interval. Key: datastore numberReadAveraged_average
Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore numberWriteAveraged_average
Read Throughput (KBps)	Amount of data read in the performance interval. Key: datastore read_average
Write Throughput (KBps)	Amount of data written to disk in the performance interval. Key: datastore write_average
Total Throughput (KBps)	Usage Average. Key: datastore usage_average
Read Latency	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore totalReadLatency_average

Table 8-2. Datastore Metrics for Storage Pods (continued)

Metric Name	Description
Write Latency	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore totalWriteLatency_average
Total 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
Total IOPS	Average number of commands issued per second during the collection interval. Key: datastore commandsAveraged_average

Table 8-3. Disk Space Metrics for Storage Pods

Metric Name	Description
Freespace	Unused space available on datastore. Key: diskspace freespace
Total used	Total space used. Key: diskspace disktotal
Capacity	Total capacity of datastore. Key: diskspace capacity
Virtual Machine used	Space used by virtual machine files. Key: diskspace used
Snapshot Space	Space used by snapshots. Key: diskspace snapshot

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 badge metrics. See definitions in [Badge Metrics](#).

Table 8-4. Network Metrics for VMware Distributed Virtual Switches

Metric Name	Description
Total Ingress Traffic	Total ingress traffic (KBps). Key: network port_statistics rx_bytes
Total Egress Traffic	Total egress traffic (KBps). Key: network port_statistics tx_bytes
Egress Unicast Packets per second	Egress unicast packets per second. Key: network port_statistics lucast_tx_pkts

Table 8-4. Network Metrics for VMware Distributed Virtual Switches (continued)

Metric Name	Description
Egress Multicast Packets per second	Egress multicast packets per second. Key: network port_statistics mcast_tx_pkts
Egress Broadcast Packets per second	Egress broadcast packets per second. Key: network port_statistics bcast_tx_pkts
Ingress Unicast Packets per second	Ingress unicast packets per second. Key: network port_statistics lucast_rx_pkts
Ingress Multicast Packets per second	Ingress multicast packets per second. Key: network port_statistics mcast_rx_pkts
Ingress Broadcast Packets per second	Ingress broadcast packets per second. Key: network port_statistics bcast_rx_pkts
Egress Dropped Packets per second	Egress dropped packets per second. Key: network port_statistics dropped_tx_pkts
Ingress Dropped Packets per second	Ingress dropped packets per second. Key: network port_statistics dropped_rx_pkts
Total Ingress Packets per second	Total ingress packets per second. Key: network port_statistics rx_pkts
Total Egress Packets per second	Total egress packets per second. Key: network port_statistics tx_pkts
Utilization	Use (KBps). Key: network port_statistics utilization
Total Dropped Packets per second	Total dropped packets per second. Key: network port_statistics dropped_pkts
Percentage of Dropped Packets	Percentage of dropped packets. Key: network port_statistics dropped_pkts_pct
Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps). Key: network port_statistics maxObserved_rx_bytes
Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps). Key: network port_statistics maxObserved_tx_bytes
Max Observed Utilization (KBps)	Max observed utilization (KBps). Key: network port_statistics maxObserved_utilization

Table 8-5. Summary Metrics for VMware Distributed Virtual Switches

Metric Name	Description
Maximum Number of Ports	Maximum number of ports. Key: summary max_num_ports
Used Number of Ports	Used number of ports. Key: summary used_num_ports
Number of Blocked Ports	Number of blocked ports. Key: summary num_blocked_ports

Table 8-6. Host Metrics for VMware Distributed Virtual Switches

Metric Name	Description
MTU Mismatch	Maximum Transmission Unit (MTU) mismatch. Key: host mtu_mismatch
Teaming Mismatch	Teaming mismatch. Key: host teaming_mismatch
Unsupported MTU	Unsupported MTU. Key: host mtu_unsupported
Unsupported VLANs	Unsupported VLANs. Key: host vlans_unsupported
Config Out Of Sync	Config Out Of Sync. Key: host config_outofsync
Number of Attached pNICs	Number of attached physical NICs. Key: host attached_pnics

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 badge metrics. See definitions in [Badge Metrics](#).

Table 8-7. Network Metrics for Distributed Virtual Port Groups

Metric Name	Description
Ingress Traffic	Ingress traffic (KBps). Key: network port_statistics rx_bytes
Egress Traffic	Egress traffic (KBps). Key: network port_statistics tx_bytes
Egress Unicast Packets per second	Egress unicast packets per second. Key: network port_statistics lucast_tx_pkts
Egress Multicast Packets per second	Egress multicast packets per second. Key: network port_statistics mcast_tx_pkts

Table 8-7. Network Metrics for Distributed Virtual Port Groups (continued)

Metric Name	Description
Egress Broadcast Packets per second	Egress broadcast packets per second. Key: network port_statistics bcast_tx_pkts
Ingress Unicast Packets per second	Ingress unicast packets per second. Key: network port_statistics lucast_rx_pkts
Ingress Multicast Packets per second	Ingress multicast packets per second. Key: network port_statistics lmcaster_rx_pkts
Ingress Broadcast Packets per second	Ingress broadcast packets per second. Key: network port_statistics bcast_rx_pkts
Egress Dropped Packets per second	Egress dropped packets per second. Key: network port_statistics dropped_tx_pkts
Ingress Dropped Packets per second	Ingress dropped packets per second. Key: network port_statistics dropped_rx_pkts
Total Ingress Packets per second	Total Ingress packets per second. Key: network port_statistics rx_pkts
Total Egress Packets per second	Total Egress packets per second. Key: network port_statistics tx_pkts
Utilization	Utilization (KBps). Key: network port_statistics utilization
Total Dropped Packets per second	Total dropped packets per second. Key: network port_statistics dropped_pkts
Percentage of Dropped Packets	Percentage of dropped packets. Key: network port_statistics dropped_pkts_pct
Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps). Key: network port_statistics maxObserved_rx_bytes
Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps). Key: network port_statistics maxObserved_tx_bytes
Max Observed Utilization (KBps)	Max observed utilization (KBps). network port_statistics maxObserved_utilization

Table 8-8. Summary Metrics for Distributed Virtual Port Groups

Metric Name	Description
Maximum Number of Ports	Maximum number of ports. Key: summary max_num_ports
Used Number of Ports	Used number of ports. Key: summary used_num_ports
Number of Blocked Ports	The number of blocked ports. Key: summary num_blocked_ports

Datastore Cluster Metrics

vRealize Operations Manager collects profile metrics for the datastore cluster resources.

Profiles Metrics for Datastore Cluster Resources

Profiles metrics provide information about the profile specific capacity.

Metric Name	Description
Profiles Capacity Remaining Profile (Average)	The capacity remaining in terms of fitting the average consumer. Key: Profiles capacityRemainingProfile_<profile uuid>
Profiles Capacity Remaining Profile (<custom profile name>)	Published for custom profiles enabled from policy on Datastore Cluster Resource. Key: Profiles capacityRemainingProfile_<profile uuid>

Capacity Allocation Metrics for Datastore Cluster Resources

Capacity allocation metrics provide information about the allotment of capacity, see [Capacity Analytics Generated Metrics](#).

Datastore Metrics

vRealize Operations Manager collects capacity, device, and summary metrics for datastore objects.

Capacity metrics can be calculated for datastore objects. See [Capacity Analytics Generated Metrics](#).

Capacity Metrics for Datastores

Capacity metrics provide information about datastore capacity.

Metric Name	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 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 being used on the datastore. Check the metric trend to identify spikes or abnormal growth. Key: capacity provisioned

Metric Name	Description
Capacity Total Capacity (GB)	<p>This metric shows the overall size of the datastore.</p> <p>Use this metric to know the total capacity of the datastore.</p> <p>Typically the size of the datastore should not be too small. VMFS datastore size has grown over the years as virtualization matures and larger virtual machines are now onboard. Ensure that the size can handle enough virtual machines to avoid datastore sprawl. A best practice is to use 5 TB for VMFS and more for vSAN.</p> <p>Key: capacity total_capacity</p>
Capacity Used Space (GB)	<p>This metric shows the amount of storage that is being used on the datastore.</p> <p>Key: capacity used_space</p>
Capacity Workload (%)	<p>Capacity workload.</p> <p>Key: capacity workload</p>
Capacity Uncommitted Space (GB)	<p>Uncommitted space in gigabytes.</p> <p>Key: capacity uncommitted</p>
Capacity Total Provisioned Consumer Space	<p>Total Provisioned Consumer Space.</p> <p>Key: capacity consumer_provisioned</p>
Capacity Used Space (%)	<p>This metric shows the amount of storage that is being used on the datastore.</p> <p>Use this metric to know the percentage of storage space being used on the datastore.</p> <p>When using this metric, verify that you have at least 20% of free storage. Less than this, and you might 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.</p> <p>Key: capacity usedSpacePct</p>

Device Metrics for Datastores

Device metrics provide information about device performance.

Metric Name	Description
Devices Bus Resets	<p>This metric shows the number of bus resets in the performance interval.</p> <p>Key: devices busResets_summation</p>
Devices Commands Aborted	<p>This metric shows the number of disk commands canceled in the performance interval.</p> <p>Key: devices commandsAborted_summation</p>
Devices Commands Issued	<p>This metric shows the number of disk commands issued in the performance interval.</p> <p>Key: devices commands_summation</p>

Metric Name	Description
Devices Read Latency (ms)	<p>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.</p> <p>Key: devices totalReadLatency_average</p>
Devices Kernel Disk Read Latency (ms)	<p>Average time spent in ESX host VM Kernel per read.</p> <p>Key: devices kernelReadLatency_average</p>
Devices Kernel Write Latency (ms)	<p>Average time spent in ESX Server VM Kernel per write.</p> <p>Key: devices kernelWriteLatency_average</p>
Devices Physical Device Read Latency (ms)	<p>Average time taken to complete a read from the physical device.</p> <p>Key: devices deviceReadLatency_average</p>
Devices Queue Write Latency (ms)	<p>Average time spent in the ESX Server VM Kernel queue per write.</p> <p>Key: devices queueWriteLatency_average</p>
Devices Physical Device Write Latency (ms)	<p>Average time taken to complete a write from the physical disk.</p> <p>Key: devices deviceWriteLatency_average</p>

Datastore Metrics for Datastores

Datastore metrics provide information about datastore use.

Metric Name	Description
Datastore Total Latency (ms)	<p>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.</p> <p>This is the average of all the VMs running in the datastore. Because it is an average, some VMs logically experience higher latency than the value shown by this metric. To see the worst latency experienced by any VM, use the Maximum VM Disk Latency metric.</p> <p>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.</p> <p>The number should be lower than the performance you expect.</p> <p>Key: datastore totalLatency_average</p>
Datastore Total Throughput (KBps)	<p>Average use in kilobytes per second.</p> <p>Key: datastore usage_average</p>

Metric Name	Description
Datastore Read Latency (ms)	<p>Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.</p> <p>Key: datastore totalReadLatency_average</p>
Datastore Write Latency (ms)	<p>Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.</p> <p>Key: datastore totalWriteLatency_average</p>
Datastore Demand	<p>Demand.</p> <p>Key: datastore demand</p>
Datastore Outstanding IO requests	<p>OIO for datastore.</p> <p>Key: datastore demand_oio</p>
Datastore Read IOPS	<p>This metric displays the average number of read commands issued per second during the collection interval.</p> <p>Use this metric when the total IOPS is higher than expected. 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.</p> <p>Key: datastore numberReadAveraged_average</p>
Datastore Write IOPS	<p>This metric displays the average number of write commands issued per second during the collection interval.</p> <p>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.</p> <p>Key: datastore numberWriteAveraged_average</p>
Datastore Read Throughput (KBps)	<p>This metric displays the amount of data read in the performance interval.</p> <p>Key: datastore read_average</p>
Datastore Write Throughput (KBps)	<p>This metric displays the amount of data written to disk in the performance interval.</p> <p>Key: datastore write_average</p>

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 Name	Description
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
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 when 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 used (GB)	Total used space in gigabytes. Key: diskspace disktotal
Diskspace Swap File Space (GB)	Swap file space in gigabytes. Key: diskspace swap

Metric Name	Description
Diskspace Other VM Space (GB)	Other virtual machine space in gigabytes. Key: diskspacelotherused
Diskspace Freespace (GB)	Unused space available on datastore. Key: diskspacelfreespace
Diskspace Capacity (GB)	Total capacity of datastore in gigabytes. Key: diskspacelcapacity
Diskspace Overhead	Amount of disk space that is overhead. Key: diskspaceloverhead

Summary Metrics for Datastores

Summary metrics provide information about overall performance.

Metric Name	Description
Summary Number of Hosts	<p>This metric shows the number of hosts that the datastore is connected to.</p> <p>Use this metric to know how many clusters the datastore is attached to.</p> <p>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.</p> <p>Key: summary total_number_hosts</p>
Summary Total Number of VMs	<p>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 is counted on each datastore.</p> <p>Use this metric to know how many VMs have at least one VMDK on a specific datastore.</p> <p>The number of VMs should be within your Concentration Risk policy.</p> <p>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.</p> <p>Key: summary total_number_vms</p>
Summary Maximum Number of VMs	<p>Maximum number of virtual machines.</p> <p>Key: summary max_number_vms</p>
Summary Workload Indicator	<p>Workload indicator.</p> <p>Key: summary workload_indicator</p>
Summary Number of Clusters	<p>This metric shows the number of clusters that the datastore is connected to.</p> <p>Key: summary total_number_clusters</p>
Summary Number of VM Templates	<p>Number of VM Templates.</p> <p>Key: Summary Number of VM Templates</p>

Template Metrics for Datastores

Metric Name	Description
Template Virtual Machine used	Space used by virtual machine files. Key: template used
Template Access Time	Last access time. Key: template accessTime

Cost Metrics for Datastores

Cost metrics provides information about the cost.

Metric Name	Description
Monthly Disk Space Base Rate	Disk space base rate for datastore displays the cost of 1 GB storage. Key: cost storageRate
Monthly Total Cost	Monthly total cost, computed by multiplying datastore capacity with monthly storage rate. Key: cost totalCost
Cost Allocation Disk Space Base Rate (Currency)	Monthly storage rate for datastore displays the cost of 1 GB storage when the overcommit ratio is set in policy. cost storageRate
Cost Allocation Monthly Datastore Allocated Cost(Currency/Month)	Monthly allocated cost as compared to the total cost of the datastore
Cost Allocation Monthly Datastore Unallocated Cost(Currency/Month)	Monthly unallocated cost as compared to the total cost of the datastore.

Reclaimable Metrics

Reclaimable metrics provide information about reclaimable resources.

Metric Name	Description
Reclaimable Orphaned Disks Disk Space (GB)	Summary of storage used by all orphaned VMDKs on the datastore. Key: reclaimable orphaned_disk diskspace
Reclaimable Orphaned Disks Potential Savings (Currency)	Potential saving after reclamation of storage by removing orphaned VMDKs from the datastore. Key: reclaimable orphaned_disk cost

Disabled Instanced Metrics

The instance metrics created for the following metrics are disabled in this version of vRealize Operations Manager . This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

Metric Name
Devices Kernel Latency (ms)
Devices Number of Running Hosts
Devices Number of Running VMs
Devices Physical Device Latency (ms)
Devices Queue Latency (ms)
Devices Queue Read Latency (ms)
Devices Read IOPS
Devices Read Latency (ms)
Devices Read Requests
Devices Read Throughput (KBps)
Devices Total IOPS
Devices Total Latency (ms)
Devices Total Throughput (KBps)
Devices Write IOPS
Devices Write Latency (ms)
Devices Write Requests
Devices Write Throughput (KBps)

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, see [Metrics and Properties Details](#).

Metric Name	Key
Capacity Data Store Capacity Contention (%)	capacity contention
Datastore I/O Demand Indicator	datastore demand_indicator
Datastore I/O Max Observed Number of Outstanding IO Operations	datastore maxObserved_OIO
Datastore I/O Max Observed Read Latency (msec)	datastore maxObserved_Read
Datastore I/O Max Observed Read Latency (msec)	datastore maxObserved_ReadLatency

Metric Name	Key
Datastore I/O Max Observed	datastore maxObserved_NumberRead
Datastore I/O Max Observed Write Latency (msec)	datastore maxObserved_Write
Datastore I/O Max Observed Write Latency (msec)	datastore maxObserved_WriteLatency
Datastore I/O Max Observed Writes per second	datastore maxObserved_NumberWrite
Datastore Demand Indicator	Demand Indicator. Key: datastore demand_indicator
Diskspace Not Shared (GB)	Unshared space in gigabytes. Key: diskspace notshared

Cluster Compute Metrics for Allocation Model

vRealize Operations Manager collects configuration, disk space, CPU use, disk, memory, network, power, and summary metrics for cluster compute resources.

Cost Metrics for Cluster Compute Resources

Cost metrics provide information about the cost.

Metric Name	Description
Cluster CPU Base Rate	Base rate for Cluster CPU calculated by dividing the monthly total cluster CPU cost by cluster CPU over-commit ratio. Key: Cost Allocation ClusterCPUBaseRate
Cluster Memory Base Rate	Cluster memory base rate calculated by dividing the monthly total cluster memory cost by cluster memory over-commit ratio. Key: Cost Allocation ClusterMemoryBaseRate
Monthly Cluster Allocated Cost	Sum of of monthly cluster CPU, Memory, and Storage costs Key: Cost Allocation MonthlyClusterAllocatedCost
Monthly Cluster Unallocated Cost	Monthly cluster unallocated cost calculated by subtracting the monthly cluster allocated cost from the monthly cluster total cost. Key: Cost Allocation MonthlyClusterUnallocatedCost
Monthly Storage Rate	Datastore base rate is calculated by dividing Storage base rate based on utilization by over commit ratio. Key: Cost Allocation Monthly Storage Rate

Virtual Machine Metrics for Allocation Model

vRealize Operations Manager collects configuration, disk space, CPU use, disk, memory, network, power, and summary metrics for virtual machine resources.

Cost Metrics for Virtual Machines

Cost metrics provide information about the cost.

Metric Name	Description
MTD VM CPU Cost	Month to date virtual machine CPU cost. Key: Cost Allocation MTD VM CPU Cost
MTD VM Memory Cost	Month to date virtual machine memory cost. Key: Cost Allocation MTD VM Memory Cost
MTD VM Storage Cost	Month to date storage cost of the virtual machine. Key: Cost Allocation MTD VM Storage Cost
MTD VM Total Cost	Addition of CPU ,Memory ,Storage, and Direct cost. Key: Cost Allocation MTD VM Total Cost

Metrics for Namespace

vRealize Operations Manager collects metrics for Namespace through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

Table 8-9. Metrics for Namespace

Metric Key	Localized Name	Description
cpu usagemhz_average	CPU Usage	Average CPU usage in MHZ.
cpu demandmhz	CPU Demand	Demand(MHz).
cpu capacity_contentionPct	CPU Contention	Percent of time descendant virtual machines are unable to run because they are contending for access to the physical CPU(s).
cpueffective_limit	CPU Effective limit	CPU Effective limit.
cpu reservation_used	CPU Reservation Used	CPU Reservation Used.
cpu estimated_entitlement	CPU Estimated entitlement	CPU Estimated entitlement.
cpudynamic_entitlement	CPU Dynamic entitlement	CPU Dynamic Entitlement.
cpu capacity_contention	CPU Overall CPU Contention	Overall CPU Contention (ms).
cpu capacity_demandEntitlementPct	CPU Capacity Demand Entitlement	CPU Capacity Demand Entitlement Percentage.
mem usage_average	Memory Usage	Memory currently in use as a percentage of total available memory.
mem guest_provisioned	Memory Total Capacity	Total Capacity.
mem active_average	Memory Guest Active	Amount of memory that is actively used.
mem granted_average	Memory Granted	Amount of memory available for use.
mem shared_average	Memory Shared	Amount of shared memory.

Table 8-9. Metrics for Namespace (continued)

Metric Key	Localized Name	Description
mem overhead_average	Memory VM Overhead	Memory overhead reported by host.
mem consumed_average	Memory Consumed	Amount of host memory consumed by the virtual machine for guest memory.
mem host_contentionPct	Memory Contention	Machine Contention Percentage.
mem guest_usage	Memory Guest Usage	Guest Memory Entitlement.
mem guest_demand	Memory Guest Demand	Guest Memory Entitlement.
mem reservation_used	Memory Reservation Used	Memory Reservation Used.
mem effective_limit	Memory Effective limit	Memory Effective limit.
mem swpinRate_average	Memory Swap In Rate	Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance.
mem swapoutRate_average	Memory Swap Out Rate	Rate at which memory is being swapped from active memory to disk during the current interval.
mem vmmemctl_average	Memory Balloon	Amount of memory currently used by the virtual machine memory control.
mem zero_average	Memory Zero	Amount of memory that is all 0.
mem swapped_average	Memory Swapped	Amount of unreserved memory.
mem zipped_latest	Memory Zipped	N/A
mem compressionRate_average	Memory Compression Rate	N/A
mem decompressionRate_average	Memory Decompression Rate	N/A
mem swpin_average	Memory Swap In	Amount of memory swapped in.
mem swapout_average	Memory Swap Out	Amount of memory swapped out.
mem swapused_average	Memory Swap Used	Amount of memory used for swap space.
mem host_contention	Memory Contention	Machine Contention.
mem dynamic_entitlement	Memory Dynamic Entitlement	Memory Dynamic Entitlement.
diskspace total_usage	Disk Space Utilization	Storage space utilized on connected vSphere Datastores.
summary configStatus	Summary Config Status	Workload Management Configuration Status.
summary total_number_pods	Summary Number of Pods	Number of Pods.
summary numberKubernetesClusters	Summary Number of Kubernetes clusters	Number of Kubernetes clusters.

Table 8-9. Metrics for Namespace (continued)

Metric Key	Localized Name	Description
summary number_running_vms	Summary Number of Running VMs	Number of Running VMs.
summary total_number_vms	Summary Total Number of VMs	Total Number of VMs.
summary iowait	Summary IO Wait	IO Wait.

Metrics for Tanzu Kubernetes cluster

vRealize Operations Manager collects metrics for Tanzu Kubernetes cluster through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

Table 8-10. Metrics for Tanzu Kubernetes clusters

Metric Key	Localized Name	Description
cpu usagemhz_average	CPU Usage	Average CPU usage in MHZ
cpu demandmhz	CPU Demand	Demand(MHz)
cpu capacity_contentionPct	CPU Contention	Percent of time descendant virtual machines are unable to run because they are contending for access to the physical CPU(s).
cpueffective_limit	CPU Effective limit	CPU Effective limit
cpu reservation_used	CPU Reservation Used	CPU Reservation Used
cpu estimated_entitlement	CPU Estimated entitlement	CPU Estimated entitlement
cpu dynamic_entitlement	CPU Dynamic entitlement	CPU Dynamic Entitlement
cpu capacity_contention	CPU Overall CPU Contention	Overall CPU Contention (ms)
cpu capacity_demandEntitlementPct	CPU Capacity Demand Entitlement	CPU Capacity Demand Entitlement Percentage
mem usage_average	Memory Usage	Memory currently in use as a percentage of total available memory
mem guest_provisioned	Memory Total Capacity	Total Capacity
mem active_average	Memory Guest Active	Amount of memory that is actively used
mem granted_average	Memory Granted	Amount of memory available for use
mem shared_average	Memory Shared	Amount of shared memory
mem overhead_average	Memory VM Overhead	Memory overhead reported by host
mem consumed_average	Memory Consumed	Amount of host memory consumed by the virtual machine for guest memory
mem host_contentionPct	Memory Contention	Machine Contention Percentage

Table 8-10. Metrics for Tanzu Kubernetes clusters (continued)

Metric Key	Localized Name	Description
mem guest_usage	Memory Guest Usage	Guest Memory Entitlement
mem guest_demand	Memory Guest Demand	Guest Memory Entitlement
mem reservation_used	Memory Reservation Used	Memory Reservation Used
mem effective_limit	Memory Effective limit	Memory Effective limit
mem swpinRate_average	Memory Swap In Rate	Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance.
mem swapoutRate_average	Memory Swap Out Rate	Rate at which memory is being swapped from active memory to disk during the current interval
mem vmmemctl_average	Memory Balloon	Amount of memory currently used by the virtual machine memory control
mem zero_average	Memory Zero	Amount of memory that is all 0
mem swapped_average	Memory Swapped	Amount of unreserved memory
mem zipped_latest	Memory Zipped	N/A
mem compressionRate_average	Memory Compression Rate	N/A
mem decompressionRate_average	Memory Decompression Rate	N/A
mem swpin_average	Memory Swap In	Amount of memory swapped in
mem swapout_average	Memory Swap Out	Amount of memory swapped out
mem swapused_average	Memory Swap Used	Amount of memory used for swap space
mem host_contention	Memory Contention	Machine Contention
mem dynamic_entitlement	Memory Dynamic Entitlement	Memory Dynamic Entitlement
summary number_running_vms	Summary Number of Running VMs	Number of Running VMs
summary total_number_vms	Summary Total Number of VMs	Total Number of VMs
summary iowait	Summary IO Wait	IO Wait

Metrics for vSphere Pods

vRealize Operations Manager collects metrics for vSphere Pods through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

Table 8-11. Metrics for vSphere Pods

Metric Key	Metric Name	Description
config hardware num_Cpu	Configuration Hardware Number of CPUs	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.
config hardware disk_Space	Configuration Hardware Disk Space	Disk space metrics
config hardware thin_Enabled	Configuration Hardware Thin Provisioned Disk	Thin Provisioned Disk
config cpuAllocation slotSize	Configuration CPU Resource Allocation HA Slot Size	vSphere HA Slot Size for CPU
config memoryAllocation slotSize	Configuration Memory Resource Allocation HA Slot Size	vSphere HA Slot Size for Memory
cpu usage_average	CPU Usage	CPU Usage divided by VM CPU Configuration in MHz
cpu usagemhz_average	CPU Usage	Amount of actively used virtual CPU. This is the host's view of the CPU usage, not the guest operating system view.
cpu usagemhz_average_mtd	CPU Usage average MTD	Month to date average CPU usage in MHZ
cpu readyPct	CPU Ready	Percentage of CPU the VM is ready to run, but unable due to ESXi has no ready physical core to run it. High Ready value impacts VM performance
cpu capacity_contentionPct	CPU Contention	Percentage of time VM is not getting the CPU resource it demanded. Impacted by Ready, Co-Stop, Hyper Threading and Power Management
cpu corecount_provisioned	CPU Provisioned vCPU(s)	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.
cpu vm_capacity_provisioned	CPU Total Capacity	Configured Capacity in MHz, based on nominal (static) frequency of the CPU
cpu demandmhz	CPU Demand	The amount of CPU resources virtual machine would use if there were no CPU contention or CPU limit.
cpu demandPct	CPU Demand (%)	The percentage of CPU resources virtual machine would use if there were no CPU contention or CPU limit.
cpu reservation_used	CPU Reservation Used	CPU Reserved for the VM. It's guaranteed to be available when the VM demands it.
cpu effective_limit	CPU Effective limit	Limit placed on the VM by vSphere. Avoid using limit as it impacts VM performance
cpu iowaitPct	CPU IO Wait	Percentage of time VM CPU is waiting for IO. Formula is Wait - Idle - Swap Wait. High value indicates slow storage subsystem

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
cpu swapwaitPct	CPU Swap wait	Percentage of time CPU is waiting on data swap-in. Mapped to vCenter CPU Swap wait
cpu costopPct	CPU Co-stop (%)	Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints. VM with less vCPU have lower co-stop value.
cpu system_summation	CPU System	CPU time spent on system processes
cpu wait_summation	CPU Wait	Total CPU time spent in wait state
cpu ready_summation	CPU Ready	CPU time spent on ready state
cpu used_summation	CPU Used	CPU time that is used
cpu iowait	CPU IO Wait	IO Wait
cpu wait	CPU Total Wait	CPU time spent on idle state
cpu capacity_demandEntitlementPct	CPU Capacity Demand Entitlement	CPU Capacity Demand Entitlement Percentage
cpu host_demand_for_aggregation	CPU Host Demand For Aggregation	Host demand for aggregation
cpu dynamic_entitlement	CPU Dynamic entitlement	CPU Dynamic entitlement
cpu capacity_contention	CPU Overall CPU Contention	Overall CPU Contention (ms)
cpu estimated_entitlement	CPU Estimated entitlement	CPU Estimated entitlement
cpu idlePct	CPU Idle	% CPU time that is idle
cpu waitPct	CPU Wait	% Total CPU time spent in wait state
cpu systemSummationPct	CPU System	% CPU time spent on system processes
cpu demandOverLimit	CPU Demand Over Limit	Amount of CPU Demand that is over the configured CPU Limit
cpu demandOverCapacity	CPU Demand Over Capacity	Amount of CPU Demand that is over the configured CPU Capacity
cpu perCpuCoStopPct	CPU Normalized Co-stop	Percentage of co-stop time, normalized across all vCPUs
cpu swapwait_summation	CPU Swap Wait	Amount of time waiting on swap.
cpu costop_summation	CPU Co-stop	Time the VM is ready to run, but is unable to due to co-scheduling constraints.
cpu idle_summation	CPU Idle	CPU time that is idle.
cpu latency_average	CPU Latency	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs.

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
cpu maxlimited_summation	CPU Max Limited	Time the VM is ready to run, but is not run due to maxing out its CPU limit setting.
cpu overlap_summation	CPU Overlap	Time the VM was interrupted to perform system services on behalf of that VM or other VMs.
cpu run_summation	CPU Run	Time the VM is scheduled to run.
cpulentitlement_latest	CPU Entitlement Latest	Entitlement Latest.
cpu demandEntitlementRatio_latest	CPU Demand-to-entitlement Ratio	CPU resource entitlement to CPU demand ratio (in percents)
cpu readiness_average	CPU Readiness	Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.
rescpu actav1_latest	CPU Utilization for Resources CPU Active (1 min. average)	The average active time for the CPU over the past minute
rescpu actav5_latestswapiRate_average	CPU Utilization for Resources CPU Active (5 min. average)	The average active time for the CPU over the past five minutes.
rescpu actav5_latest	CPU Utilization for Resources CPU Active (5 min. average)	The average active time for the CPU over the past five minutes
rescpu actav15_latest	CPU Utilization for Resources CPU Active (15 min. average)	The average active time for the CPU over the past fifteen minutes
rescpu actpk1_latest	CPU Utilization for Resources CPU Active (1 min. peak)	The peak active time for the CPU over the past minute
rescpu actpk5_latest	CPU Utilization for Resources CPU Active (5 min. peak)	The peak active time for the CPU over the past five minutes
rescpu actpk15_latest	CPU Utilization for Resources CPU Active (15 min. peak)	The peak active time for the CPU over the past fifteen minutes
rescpu runav1_latest	CPU Utilization for Resources CPU Running (1 min. average)	The average runtime for the CPU over the past minute
rescpu runav5_latest	CPU Utilization for Resources CPU Running (5 min. average)	The average runtime for the CPU over the past five minutes
rescpu runav15_latest	CPU Utilization for Resources CPU Running (15 min. average)	The average runtime for the CPU over the past fifteen minutes
rescpu runpk1_latest	CPU Utilization for Resources CPU Running (1 min. peak)	The peak active time for the CPU over the past minute
rescpu runpk5_latest	CPU Utilization for Resources CPU Running (5 min. peak)	The peak active time for the CPU over the past five minutes
rescpu runpk15_latest	CPU Utilization for Resources CPU Running (15 min. peak)	The peak active time for the CPU over the past fifteen minutes
rescpu maxLimited1_latest	CPU Utilization for Resources CPU Throttled (1 min. average)	The scheduling limit over the past minute

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
rescpu maxLimited5_latest	CPU Utilization for Resources CPU Throttled (5 min. average)	The scheduling limit over the past five minutes
rescpu maxLimited15_latest	CPU Utilization for Resources CPU Throttled (15 min. average)	The scheduling limit over the past fifteen minutes
rescpu sampleCount_latest	CPU Utilization for Resources Group CPU Sample Count	The sample CPU count
rescpu samplePeriod_latest	CPU Utilization for Resources Group CPU Sample Period	The sample period
mem usage_average	Memory Usage	Memory currently in use as a percentage of total available memory
mem balloonPct	Memory Balloon	Percentage of guest physical memory that is currently claimed from the virtual machine through ballooning. This is the percentage of guest physical memory that has been allocated and pinned by the balloon driver. Balloon does not necessarily mean the VM performance is affected.
mem swapped_average	Memory Swapped	Amount of unreserved memory
mem consumed_average	Memory Consumed	Amount of ESXi Host memory mapped/ consumed by the virtual machine for guest memory
mem consumed_average_mtd	Memory Consumed average MTD	average MTD Amount of host memory consumed by the virtual machine for guest memory
mem consumedPct	Memory Consumed (%)	Amount of host memory consumed by the virtual machine for guest memory. Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used.
mem overhead_average	Memory Overhead	Amount of overhead memory used by ESXi to run the Virtual Machine.
mem host_contentionPct	Memory Contention	Percentage of time the VM has contended for memory.
mem guest_provisioned	Memory Total Capacity	Memory resources allocated to the Virtual Machine
mem guest_usage	Memory Guest Usage	Guest Memory Entitlement
mem guest_demand	Memory Guest Demand	Guest Memory Entitlement
mem host_demand	Memory Host Demand	Memory Demand in KB
mem reservation_used	Memory Reservation Used	Memory Reservation Used
mem effective_limit	Memory Effective limit	Memory Effective limit

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
mem vmMemoryDemand	Memory Utilization	Amount of memory utilized by the Virtual Machine. Reflects the guest OS memory required (for certain vSphere and VMTools versions) or Virtual Machine consumption
mem nonzero_active	Memory Non Zero Active	Non Zero Active Memory
mem swpinRate_average	Memory Swap In Rate	Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance.
mem swapoutRate_average	Memory Swap Out Rate	Rate at which memory is being swapped from active memory to disk during the current interval.
mem compressed_average	Memory Compressed	Percentage of total memory that has been compressed by vSphere. If and only if the page is accessed by the Guest OS, will performance be affected.
mem overheadMax_average	Memory Overhead Max	N/A
mem vmmemctl_average	Memory Balloon	Amount of memory currently used by the virtual machine memory control
mem active_average	Memory Guest Active	Amount of memory that is actively used
mem granted_average	Memory Granted	Amount of memory available for use
mem shared_average	Memory Shared	Amount of shared memory
mem zero_average	Memory Zero	Amount of memory that is all 0
mem swaptarget_average	Memory Swap Target	Amount of memory that can be swapped
mem swpin_average	Memory Swap In	Amount of memory swapped in
mem swapout_average	Memory Swap Out	Amount of memory swapped out
mem vmmemctltarget_average	Memory Balloon Target	Amount of memory that can be used by the virtual machine memory control
mem host_dynamic_entitlement	Memory Host Dynamic Entitlement	Mem Machine Dynamic Entitlement
mem host_active	Memory Host Active	Machine Active
mem host_usage	Memory Host Usage	Machine Usage
mem host_contention	Memory Contention	Machine Contention
mem guest_activePct	Memory Guest Active Memory	Guest active memory as percentage of configured
mem guest_dynamic_entitlement	Memory Guest Dynamic Entitlement	Guest Memory Dynamic Entitlement

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
mem host_demand_reservation	Memory Host Demand with Reservation	Memory Demand with Reservation considered in KB
mem host_nonpageable_estimate	Memory Guest Non Pageable Memory	Guest Non Pageable Memory Estimates
mem guest_nonpageable_estimate	Memory Host Non Pageable Memory	Guest Non Pageable Memory Estimates
mem estimated_entitlement	Memory Estimated entitlement	Memory Estimated entitlement
mem host_demand_for_aggregation	Memory Host Demand For Aggregation	Host demand for aggregation
mem demandOverLimit	Memory Demand Over Limit	Amount of Memory Demand that is over the configured Memory Limit
mem demandOverCapacity	Memory Demand Over Capacity	Amount of Memory Demand that is over the configured Memory Capacity
mem activewrite_average	Memory Active Write	N/A
mem compressionRate_average	Memory Compression Rate	N/A
mem decompressionRate_average	Memory Decompression Rate	N/A
mem zipSaved_latest	Memory Zip Saved	N/A
mem zipped_latest	Memory Zipped	N/A
mem entitlement_average	Memory Entitlement	Amount of host physical memory the VM is entitled to, as determined by the ESX schedule.
mem latency_average	Memory Latency	Percentage of time the VM is waiting to access swapped or compressed memory.
mem capacity.contention_average	Memory Capacity Contention	Capacity Contention.
mem ISwapInRate_average	Memory Swap In Rate from Host Cache	Rate at which memory is being swapped from host cache into active memory.
mem ISwapOutRate_average	Memory Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory.
mem ISwapUsed_average	Memory Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache.
mem overheadTouched_average	Memory Overhead Touched	Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM.
net usage_average	Network Usage Rate	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine
net transmitted_average	Network Data Transmit Rate	Average amount of data transmitted per second

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
net received_average	Network Data Receive Rate	Average amount of data received per second
net droppedTx_summation	Network Transmitted Packets Dropped	Number of outgoing packets dropped in the performance interval. Investigate if the number is not 0
net droppedPct	Network Packets Dropped (%)	Percentage of packets dropped
net dropped	Network Packets Dropped	Number of packets dropped in the performance interval
net broadcastTx_summation	Network Broadcast Packets Transmitted	Total number of broadcast packets transmitted. Investigate further if this number is high
net multicastTx_summation	Network Multicast Packets Transmitted	Number of multicast packets transmitted. Investigate further if this number is high
net idle	Network idle	N/A
net usage_capacity	Network I/O Usage Capacity	I/O Usage Capacity
net maxObserved_KBps	Network Max Observed Throughput	Max observed rate of network throughput
net maxObserved_Tx_KBps	Network Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput
net maxObserved_Rx_KBps	Network Max Observed Received Throughput	Max observed received rate of network throughput
net packetsRx_summation	Network Packets Received	Number of packets received in the performance interval
net packetsTx_summation	Network Packets Transmitted	Number of packets transmitted in the performance interval
net demand	Network Demand	N/A
net packetsRxPerSec	Network Packets Received per second	Number of packets received in the performance interval
net packetsTxPerSec	Network Packets Transmitted per second	Number of packets transmitted in the performance interval
net packetsPerSec	Network Packets per second	Number of packets transmitted and received per second
net droppedRx_summation	Network Received Packets Dropped	Number of received packets dropped in the performance interval
net broadcastRx_summation	Network Broadcast Packets Received	Number of broadcast packets received during the sampling interval
net multicastRx_summation	Network Multicast Packets Received	Number of multicast packets received
net bytesRx_average	Network bytesRx	Average amount of data received per second

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
net bytesTx_average	Network bytesTx	Average amount of data transmitted per second
net host_transmitted_average	Network VM to Host Data Transmit Rate	Average amount of data transmitted per second between VM and host
net host_received_average	Network VM to Host Data Receive Rate	Average amount of data received per second between VM and host
net host_usage_average	Network VM to Host Usage Rate	The sum of the data transmitted and received for all the NIC instances between VM and host
net host_maxObserved_Tx_KBps	Network VM to Host Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput between VM and host
net host_maxObserved_Rx_KBps	Network VM to Host Max Observed Received Throughput	Max observed received rate of network throughput between VM and host
net host_maxObserved_KBps	Network VM to Host Max Observed Throughput	Max observed rate of network throughput between VM and host
net transmit_demand_average	Network Data Transmit Demand Rate	Data Transmit Demand Rate
net receive_demand_average	Network Data Receive Demand Rate	Data Receive Demand Rate
disk usage_average	Physical Disk Total Throughput	Amount of data read from/written to storage in a second. This is averaged over the reporting period
disk read_average	Physical Disk Read Throughput	Amount of data read from storage in a second. This is averaged over the reporting period
disk write_average	Physical Disk Write Throughput	Amount of data written to storage in a second. This is averaged over the reporting period
disk usage_capacity	Physical Disk I/O Usage Capacity	I/O Usage Capacity
disk busResets_summation	Physical Disk Bus Resets	The number of bus resets in the performance interval
disk commandsAborted_summation	Physical Disk Commands Aborted	The number of disk commands stopped in the performance interval
disk diskoio	Physical Disk Number of Outstanding IO Operations	Number of Outstanding IO Operations
disk diskqueued	Physical Disk Queued Operations	Queued Operations
disk diskdemand	Physical Disk Demand	Demand
disk sum_queued_oio	Physical Disk Total Queued Outstanding operations	Sum of Queued Operation and Outstanding Operations.

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
disk max_observed	Physical Disk Max Observed OIO	Max Observed IO for a disk.
disk numberReadAveraged_average	Physical Disk Read IOPS	Number of read operations per second. This is averaged over the reporting period.
disk numberWriteAveraged_average	Physical Disk Write IOPS	Number of write operations per second. This is averaged over the reporting period.
disk maxTotalLatency_latest	Physical Disk Highest Latency	Highest Latency.
disk scsiReservationConflicts_summation	Physical Disk SCSI Reservation Conflicts	SCSI Reservation Conflicts.
disk totalReadLatency_average	Physical Disk Read Latency	Average amount of time for a read operation by the storage adapter.
disk totalWriteLatency_average	Physical Disk Write Latency	Average amount of time for a write operation by the storage adapter.
disk totalLatency_average	Physical Disk Total Latency	Total Latency.
sys poweredOn	System Powered ON	1 if the VM is connected (available for management) and powered on, otherwise 0.
sys osUptime_latest	System OS Uptime	Total time elapsed, in seconds, since last operating system boot-up
sys uptime_latest	System Uptime	Number of seconds since system startup
sys heartbeat_summation	System Heartbeat	Number of heart beats from the virtual machine in the defined interval
sys vmotionEnabled	System vMotion Enabled	1 if vMotion enabled, 0 if not enabled
sys productString	System Product String	VMware product string
sys heartbeat_latest	System Heartbeat Latest	Number of heartbeats issued per virtual machine during the interval
summary running	Summary Running	Running
summary desktop_status	Summary Desktop Status	Horizon View Desktop Status
summary poweredOff	Summary Reclaimable Powered Off	Powered Off = 1. Not powered off = 0
summary idle	Summary Reclaimable Idle	Idle = 1. Not idle = 0
summary oversized	Summary Is Oversized	Oversized = 1. Not oversized = 0
summary undersized	Summary Is Undersized	Is Undersized
summary snapshotSpace	Summary Reclaimable Snapshot Space	Reclaimable Snapshot Space
summary oversized vcpus	Summary Oversized Virtual CPUs	Virtual CPUs

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
summary oversized memory	Summary Oversized Memory	Memory
summary undersized vcpus	Summary Undersized Virtual CPUs	Virtual CPUs
summary undersized memory	Summary Undersized Memory	Memory
summary metering value	Summary Metering Total price	Total price of the resource(Sum of all price components)
summary metering storage	Summary Metering Storage price	Price of Storage related components of the resource
summary metering memory	Summary Metering Memory price	Price of Memory related components of the resource
summary metering cpu	Summary Metering CPU price	Price of CPU related components of the resource
summary metering additional	Summary Metering Additional price	Price of additional components of the resource
summary metering partialPrice	Summary Metering Partial price	Shows whether the calculated price is partial for the resource
summary workload_indicator	Summary Workload Indicator	Workload Indicator
summary cpu_shares	Summary CPU Shares	CPU Shares
summary mem_shares	Summary Memory Shares	Memory Shares
summary number_datastore	Summary Number of Datastores	Number of Datastores
summary number_network	Summary Number of Networks	Number of Networks
guestfilesystem capacity	Guest File System Partition Capacity	Disk space capacity on guest file system partition.
guestfilesystem percentage	Guest File System Partition Utilization (%)	Guest file system partition space utilization in percentage
guestfilesystem usage	Guest File System Partition Utilization	Guest file system partition space utilization
guestfilesystem capacity_total	Guest File System Total Capacity	Disk space capacity on guest file system
guestfilesystem percentage_total	Guest File System Utilization (%)	Guest file system disk space utilization in percentage
guestfilesystem usage_total	Guest File System Utilization	Guest file system disk space utilization
guestfilesystem freespace	Guest File System Guest File System Free	Total free space on guest file system
guestfilesystem capacity_property	Guest File System Guest File System Capacity Property	Total capacity of guest file system as a property
guestfilesystem freespace_total	Guest File System Total Guest File System Free	Total free space on guest file system

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
guestfilesystem capacity_property_total	Guest File System Total Capacity Property	Total capacity of guest file system as a property
guest mem.free_latest	Guest Free Memory	Free Memory
guest mem.needed_latest	Guest Needed Memory	Needed Memory
guest mem.physUsable_latest	Guest Physically Usable Memory	Physically Usable Memory
guest page.inRate_latest	Guest Page In Rate per second	Page In Rate per second
guest page.size_latest	Guest Page Size	Page Size
guest swap.spaceRemaining_latest	Guest Remaining Swap Space	Remaining Swap Space
guest cpu_queue	Guest CPU Queue	The number of ready threads queuing in the CPU. Linux includes threads in running state. A number greater than 2 for prolong period indicates CPU core bottleneck.
guest disk_queue	Guest Disk Queue	The number of outstanding requests + IO currently in progress.
guest contextSwapRate_latest	Guest Context Swap Rate per second	Context Swap Rate per second
guest hugePage.size_latest	Guest Huge Page Size	Huge Page Size
guest hugePage.total_latest	Guest Total Huge Pages	Total Huge Pages
guest mem.activeFileCache_latest	Guest Active File Cache Memory	Active File Cache Memory
guest page.outRate_latest	Guest Page Out Rate per second	Page Out Rate per second
guest disk_queue_latest	Guest Disk Queue Latest	The number of outstanding requests + IO currently in progress.
virtualDisk numberReadAveraged_average	Virtual Disk Read IOPS	Number of read operations per second. This is averaged over the reporting period
virtualDisk numberWriteAveraged_average	Virtual Disk Write IOPS	Number of write operations per second. This is averaged over the reporting period
virtualDisk read_average	Virtual Disk Read Throughput	Amount of data read from storage in a second. This is averaged over the reporting period
virtualDisk totalReadLatency_average	Virtual Disk Read Latency	Average amount of time for a read operation by the storage adapter.
virtualDisk totalWriteLatency_average	Virtual Disk Write Latency	Average amount of time for a write operation by the storage adapter.
virtualDisk write_average	Virtual Disk Write Throughput	Amount of data written to storage in a second. This is averaged over the reporting period

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
virtualDisk usage	Virtual Disk Total Throughput	Amount of data read from/written to storage in a second. This is averaged over the reporting period
virtualDisk totalLatency	Virtual Disk Total Latency	Total Latency
virtualDisk commandsAveraged_average	Virtual Disk Total IOPS	Number of read/write operations per second. This is averaged over the reporting period
virtualDisk vDiskOIO	Virtual Disk Outstanding IO requests	OIO for datastore.
virtualDisk actualUsage	Virtual Disk Used Disk Space	Virtual Disk space usage
virtualDisk busResets_summation	Virtual Disk Bus Resets	The number of bus resets in the performance interval
virtualDisk commandsAborted_summation	Virtual Disk Commands Aborted	The number of disk commands stopped in the performance interval
virtualDisk readLoadMetric_latest	Virtual Disk Read Load	Storage DRS virtual disk metric read load
virtualDisk readOIO_latest	Virtual Disk Outstanding Read Requests	Average number of outstanding read requests to the virtual disk
virtualDisk writeLoadMetric_latest	Virtual Disk Write Load	Storage DRS virtual disk write load
virtualDisk writeOIO_latest	Virtual Disk Outstanding Write Requests	Average number of outstanding write requests to the virtual disk
virtualDisk smallSeeks_latest	Virtual Disk Number of Small Seeks	Small Seeks
virtualDisk mediumSeeks_latest	Virtual Disk Number of Medium Seeks	Medium Seeks
virtualDisk largeSeeks_latest	Virtual Disk Number of Large Seeks	Large Seeks
virtualDisk readLatencyUS_latest	Virtual Disk Read Latency (microseconds)	Read latency in microseconds
virtualDisk writeLatencyUS_latest	Virtual Disk Write Latency (microseconds)	Write Latency in microseconds
virtualDisk readIOSize_latest	Virtual Disk Average Read request size	Read IO size
virtualDisk writeIOSize_latest	Virtual Disk Average Write request size	Write IO size
diskspace pod_used	Disk Space Pod used	Space used by Pod files
diskspace provisionedSpace	Disk Space Provisioned Space for Pod	Provisioned space for Pod. In thin provisioned, it is the full space allocated (which may not be used yet).
diskspace notshared	Disk Space Not Shared	Space used by VM that is not shared with other VM

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
diskspace activeNotShared	Disk Space Active not shared	Unshared disk space used by VMs excluding snapshot
diskspace perDsUsed	Disk Space Pod used	Space used by all files of the Pod on the datastore (disks, snapshots, configs, logs, etc).
diskspace total_usage	Disk Space Utilization	Total disk space used on all datastores visible to this object
diskspace total_capacity	Disk Space Total Capacity	Total disk space on all datastores visible to this object
diskspace diskused	Disk Space Virtual Disk Used	Space used by virtual disks
diskspace snapshot	Disk Space Snapshot Space	Space used by snapshots
diskspace shared	Disk Space Shared Used	Shared space used
diskspace provisioned	Disk Space Provisioned Space	Provisioned space
diskspace snapshot used	Disk Space Snapshot Pod used	Disk space used by the Pod snapshot files. This is the space that can be potentially reclaimed if the snapshot is removed.
diskspace snapshot accessTime	Disk Space Snapshot Access Time	The date and time the snapshot was taken.
storage totalReadLatency_average	Storage Read Latency	Average amount of time for a read operation.
storage totalWriteLatency_average	Storage Write Latency	Average amount of time for a write operation.
storage read_average	Storage Read Rate	Read throughput rate
storage write_average	Storage Write Rate	Write throughput rate
storage usage_average	Storage Total Usage	Total throughput rate
storage numberReadAveraged_average	Storage Reads per second	Average number of read commands issued per second during the collection interval
storage numberWriteAveraged_average	Storage Writes per second	Average number of write commands issued per second during the collection interval
storage commandsAveraged_average	Storage Commands per second	Average number of commands issued per second during the collection interval
storage totalLatency_average	Storage Total Latency	Total latency
storage demandKBps	Storage Demand	N/A
storage contention	Storage Contention percentage	N/A
cost monthlyTotalCost	Cost MTD Total Cost	Month To Date Cost of Virtual Machine
cost monthlyProjectedCost	Cost Monthly Projected Total Cost	Virtual Machine cost projected for full month

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
cost compTotalCost	Cost MTD Compute Total Cost	Month to Date Total Compute Cost (Including CPU and Memory) of Virtual Machine
cost directCost	Cost Monthly Direct Cost	Monthly Direct Cost (comprising of OS Labor, VI Labor and any windows desktop instance license) of Virtual Machine
cost cpuCost	Cost MTD CPU Cost	Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost.
cost memoryCost	Cost MTD Memory Cost	Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost.
cost storageCost	Cost MTD Disk Space Cost	Month to Date Disk Space Cost of Virtual Machine
cost reclaimableCost	Cost Potential Savings	Potential Savings
cost osLaborTotalCost	Cost Monthly OS Labor Cost	Operating System Labor Cost of Virtual Machine for full month
cost viLaborTotalCost	Cost Monthly VI Labor Cost	Monthly VI Labor Cost
cost effectiveTotalCost	Cost MTD Effective Total Cost	Month to Date Cost of Virtual Machine considering the allocation and demand model
cost effectiveProjectedTotalCost	Cost Monthly Effective Projected Total Cost	Virtual Machine cost projected for full month considering the allocation and demand model
cost allocation allocationBasedCpuMTDCost	Cost Allocation MTD CPU Cost	Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost.
cost allocation allocationBasedMemoryMTDCost	Cost Allocation MTD Memory Cost	Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost.
cost allocation allocationBasedStorageMTDCost	Cost Allocation MTD Disk Space Cost	Month to Date Disk Space Cost of Virtual Machine
cost allocation allocationBasedTotalMTDCost	Cost Allocation MTD Total Cost	Month To Date Cost of Virtual Machine
cost allocation allocationBasedTotalCost	Cost Allocation Monthly Projected Total Cost	Virtual Machine cost projected for full month
datastore demand_oio	Datastore Outstanding IO requests	Amount of IO waiting in the queue to be executed. High IO, coupled with high latency, impacts performance.
datastore numberReadAveraged_average	Datastore Read IOPS	Number of read operations per second. This is averaged over the reporting period.
datastore numberWriteAveraged_average	Datastore Write IOPS	Number of write operations per second. This is averaged over the reporting period.

Table 8-11. Metrics for vSphere Pods (continued)

Metric Key	Metric Name	Description
datastore read_average	Datastore Read Throughput	Amount of data read from storage in a second. This is averaged over the reporting period.
datastore totalReadLatency_average	Datastore Read Latency	Average amount of time for a read operation at the datastore level. It's an average of all the VMs in the datastore.
datastore totalWriteLatency_average	Datastore Write Latency	Average amount of time for a write operation by the storage adapter.
datastore write_average	Datastore Write Throughput	Amount of data written from storage in a second. This is averaged over the reporting period.
datastore totalLatency_average	Datastore Total Latency	Normalized Latency, taking into account the read/write ratio.
datastore usage_average	Datastore Total Throughput	Amount of data read from/written to storage in a second. This is averaged over the reporting period.
datastore commandsAveraged_average	Datastore Total IOPS	Number of read/write operations per second. This is averaged over the reporting period.
datastore used	Datastore Used Space	Used Space.
datastore demand	Datastore Demand	Max of datastore "Reads Per Sec", "Writes Per Sec", "Read Rate", "Write Rate", "OIO Per Sec" percentages.
datastore maxTotalLatency_latest	Datastore Highest Latency	Highest Latency.
datastore totalLatency_max	Datastore Total Latency Max	Total Latency Max (ms).
datastore maxObserved_NumberRead	Datastore Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore maxObserved_Read	Datastore Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore maxObserved_NumberWrite	Datastore Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore maxObserved_Write	Datastore Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore maxObserved_OIO	Datastore Max Observed Number of Outstanding IO Operations	N/A

Operating System Metrics

Metrics are collected for Linux and Windows operating systems. The metrics are collected after the vRealize Application Remote Collector agent is deployed on the VM.

Linux Platforms

The following metrics are collected for Linux operating systems:

Table 8-12. Metrics for Linux

Metric	Metric Category	KPI
<Instance name> Usage Idle	CPU	False
<Instance name> Usage IO-Wait	CPU	False
<Instance name> Time Active	CPU	True
<Instance name> Time Guest	CPU	False
<Instance name> Time Guest Nice	CPU	False
<Instance name> Time Idle	CPU	False
<Instance name> Time IO-Wait	CPU	False
<Instance name> Time IRQ	CPU	True
<Instance name> Time Nice	CPU	False
<Instance name> Time Soft IRQ	CPU	True
<Instance name> Time Steal	CPU	False
<Instance name> Time System	CPU	False
<Instance name> Time User	CPU	True
<Instance name> Usage Active (%)	CPU	True
<Instance name> Usage Guest (%)	CPU	False
<Instance name> Usage Guest Nice (%)	CPU	False
<Instance name> Usage IRQ (%)	CPU	True
<Instance name> Usage Nice (%)	CPU	False
<Instance name> Usage Soft IRQ (%)	CPU	True
<Instance name> Usage Steal (%)	CPU	False
<Instance name> Usage System (%)	CPU	True
<Instance name> Usage User (%)	CPU	True
IO Time	Disk	False
Read Time	Disk	False
Reads	Disk	False
Write Time	Disk	False

Table 8-12. Metrics for Linux (continued)

Metric	Metric Category	KPI
Writes	Disk	False
<Instance name> Disk Free	Disk	False
<Instance name> Disk Total	Disk	False
<Instance name> Disk Used (%)	Disk	False
Cached	Memory	False
Free	Memory	False
Inactive	Memory	False
Total	Memory	True
Used	Memory	True
Used Percent	Memory	True
Blocked	Processes	True
Dead	Processes	False
Running	Processes	False
Sleeping	Processes	False
Stopped	Processes	False
Zombies	Processes	False
Free	Swap	False
In	Swap	False
Out	Swap	False
Total	Swap	True
Used	Swap	True
Used Percent	Swap	True

Windows Platforms

The following metrics are collected for Windows operating systems:

Table 8-13. Metrics for Windows

Metric	Metric Category	KPI
Idle Time	CPU	False
Interrupt Time	CPU	False

Table 8-13. Metrics for Windows (continued)

Metric	Metric Category	KPI
Interrupts persec	CPU	True
Privileged Time	CPU	False
Processor Time	CPU	False
User Time	CPU	False
Avg. Disk Bytes Read	Disk	False
Avg. Disk sec Read	Disk	False
Avg. Disk sec Write	Disk	False
Avg. Disk Write Queue Length	Disk	False
Avg. Disk Read Queue Length	Disk	False
Disk Read Time	Disk	False
Disk Write Time	Disk	False
Free Megabytes	Disk	False
Free Space	Disk	False
Idle Time	Disk	False
Split IO persec	Disk	False
Available Bytes	Memory	True
Cache Bytes	Memory	False
Cache Faults persec	Memory	False
Committed Bytes	Memory	True
Demand Zero Faults persec	Memory	False
Page Faults persec	Memory	True
Pages persec	Memory	False
Pool Nonpaged Bytes	Memory	True
Pool Paged Bytes	Memory	False
Transition Faults persec	Memory	False
Elapsed Time	Process	False
Handle Count	Process	False
IO Read Bytes persec	Process	False

Table 8-13. Metrics for Windows (continued)

Metric	Metric Category	KPI
IO Read Operations persec	Process	False
IO Write Bytes persec	Process	False
IO Write Operations persec	Process	False
Privileged Time	Process	False
Processor Time	Process	False
Thread Count	Process	False
User Time	Process	False
Context Switches persec	System	False
Processes	System	False
Processor Queue Length	System	False
System Calls persec	System	False
System Up Time	System	False
Threads	System	False

Application Service Metrics

Metrics are collected for 20 application services.

Active Directory Metrics

Metrics are collected for the Active Directory application service.

Table 8-14. Active Directory Metrics

Metric Name	Category	KPI
Database Cache % Hit (%)	Active Directory Database	True
Database Cache Page Faults/sec	Active Directory Database	True
Database Cache Size	Active Directory Database	False
Data Lookups	Active Directory DFS Replication	False
Database Commits	Active Directory DFS Replication	True
Avg Response Time	Active Directory DFSN	True
Requests Failed	Active Directory DFSN	False
Requests Processed	Active Directory DFSN	False

Table 8-14. Active Directory Metrics (continued)

Metric Name	Category	KPI
Dynamic Update Received	Active Directory DNS	False
Dynamic Update Rejected	Active Directory DNS	False
Recursive Queries	Active Directory DNS	False
Recursive Queries Failure	Active Directory DNS	False
Secure Update Failure	Active Directory DNS	False
Total Query Received	Active Directory DNS	True
Total Response Sent	Active Directory DNS	True
Digest Authentications	Active Directory Security System-Wide Statistics	True
Kerberos Authentications	Active Directory Security System-Wide Statistics	True
NTLM Authentications	Active Directory Security System-Wide Statistics	True
Directory Services:<InstanceName> Base Searches persec	Active Directory Services	False
Directory Services:<InstanceName> Database adds persec	Active Directory Services	False
Directory Services:<InstanceName> Database deletes persec	Active Directory Services	False
Directory Services<InstanceName> Database modifies/sec	Active Directory Services	False
Directory Services<InstanceName> Database recycles/sec	Active Directory Services	False
Directory Services<InstanceName> DRA Inbound Bytes Total/sec	Active Directory Services	False
Directory Services<InstanceName> DRA Inbound Objects/sec	Active Directory Services	False
Directory Services<InstanceName> DRA Outbound Bytes Total/sec	Active Directory Services	False
Directory Services<InstanceName> DRA Outbound Objects/sec	Active Directory Services	False
Directory Services<InstanceName> DRA Pending Replication Operations	Active Directory Services	False
Directory Services<InstanceName> DRA Pending Replication Synchronizations	Active Directory Services	False
Directory Services<InstanceName> DRA Sync Requests Made	Active Directory Services	False

Table 8-14. Active Directory Metrics (continued)

Metric Name	Category	KPI
Directory Services<InstanceName> DRA Sync Requests Successful	Active Directory Services	False
Directory Services<InstanceName> DS Client Binds/sec	Active Directory Services	True
Directory Services<InstanceName> DS Directory Reads/sec	Active Directory Services	False
Directory Services<InstanceName> DS Directory Searches/sec	Active Directory Services	True
Directory Services<InstanceName> DS Server Binds/sec	Active Directory Services	True
Directory Services<InstanceName> DS Threads in Use	Active Directory Services	True
Directory Services:<InstanceName> LDAP Active Threads	Active Directory Services	False
Directory Services:<InstanceName> LDAP Client Sessions	Active Directory Services	True
Directory Services<InstanceName> LDAP Closed Connections/sec	Active Directory Services	False
Directory Services<InstanceName> LDAP New Connections/sec	Active Directory Services	True
Directory Services<InstanceName> LDAP Searches/sec	Active Directory Services	True
Directory Services<InstanceName> LDAP Successful Binds/sec	Active Directory Services	False
Directory Services<InstanceName> LDAP UDP operations/sec	Active Directory Services	False
Directory Services:<InstanceName> LDAP Writes/sec	Active Directory Services	False

No metrics are collected for the category Active Directory.

ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

Table 8-15. ActiveMQ Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Active MQ	False
Buffer Pool<InstanceName> Memory Used	Active MQ	False

Table 8-15. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
Buffer Pool<InstanceName> Total Capacity	Active MQ	False
Class Loading Loaded Class Count	Active MQ	False
Class Loading Unloaded Class Count	Active MQ	False
Class Loading Total Loaded Class Count	Active MQ	False
File Descriptor Usage Max File Descriptor Count	Active MQ	False
File Descriptor Usage Open File Descriptor Count	Active MQ	False
Garbage Collection<InstanceName> Total Collection Count	Active MQ	False
Garbage Collection<InstanceName> Total Collection Time	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Initial Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName> Peak Usage Used Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Committed Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Initial Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Maximum Memory	Active MQ	False
JVM Memory Pool<InstanceName> Usage Used Memory	Active MQ	False

Table 8-15. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
Threading Thread Count	Active MQ	False
Uptime	Active MQ	False
UTILIZATION Process CpuLoad	Active MQ	False
UTILIZATION Memory Limit	ActiveMQ Broker	True
UTILIZATION Memory Percent Usage (%)	ActiveMQ Broker	True
UTILIZATION Store Limit	ActiveMQ Broker	False
UTILIZATION Store Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION Temp Limit	ActiveMQ Broker	False
UTILIZATION Temp Percent Usage (%)	ActiveMQ Broker	False
UTILIZATION Total Consumer Count	ActiveMQ Broker	True
UTILIZATION Total Dequeue Count	ActiveMQ Broker	True
UTILIZATION Total Enqueue Count	ActiveMQ Broker	True
UTILIZATION Total Message Count	ActiveMQ Broker	True
JVM Memory Heap Memory Usage Initial Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Heap Memory Usage Used Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Committed Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Initial Memory	ActiveMQ JVM Memory Usage	False

Table 8-15. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
JVM Memory Non Heap Memory Usage Maximum Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Non Heap Memory Usage Used Memory	ActiveMQ JVM Memory Usage	False
JVM Memory Object Pending FinalizationCount	ActiveMQ JVM Memory Usage	False
UTILIZATION Process CpuLoad	ActiveMQ OS	False
UTILIZATION System Cpu Load	ActiveMQ OS	False
UTILIZATION Consumer Count	ActiveMQ Topic	True
UTILIZATION Dequeue Count	ActiveMQ Topic	True
UTILIZATION Enqueue Count	ActiveMQ Topic	True
UTILIZATION Queue Size	ActiveMQ Topic	True
UTILIZATION Producer Count	ActiveMQ Topic	False

Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

Note Metrics are collected for the Events MPM. Metrics are not collected for the other MPMs.

Table 8-16. Apache HTTPD Metrics

Metric Name	Category	KPI
UTILIZATION Busy Workers	Apache HTTPD	True
UTILIZATION Bytes Per Req	Apache HTTPD	False
UTILIZATION Bytes Per Sec	Apache HTTPD	False
UTILIZATION CPU Load	Apache HTTPD	True
UTILIZATION CPU User	Apache HTTPD	False
UTILIZATION Idle Workers	Apache HTTPD	True
UTILIZATION Request Per Sec	Apache HTTPD	True
UTILIZATION SCBoard Closing	Apache HTTPD	False

Table 8-16. Apache HTTPD Metrics (continued)

Metric Name	Category	KPI
UTILIZATION SCBoard DNS Lookup	Apache HTTPD	False
UTILIZATION SCBoard Finishing	Apache HTTPD	False
UTILIZATION SCBoard Idle Cleanup	Apache HTTPD	False
UTILIZATION SCBoard Keep Alive	Apache HTTPD	False
UTILIZATION SCBoard Logging	Apache HTTPD	False
UTILIZATION SCBoard Open	Apache HTTPD	False
UTILIZATION SCBoard Reading	Apache HTTPD	False
UTILIZATION SCBoard Sending	Apache HTTPD	False
UTILIZATION SCBoard Starting	Apache HTTPD	False
UTILIZATION SCBoard Waiting	Apache HTTPD	False
UTILIZATION Total Accesses	Apache HTTPD	False
UTILIZATION Total Bytes	Apache HTTPD	True
UTILIZATION Total Connections	Apache HTTPD	False
UTILIZATION Uptime	Apache HTTPD	True
UTILIZATION Asynchronous Closing Connections	Apache HTTPD	False
UTILIZATION Asynchronous Keep Alive Connections	Apache HTTPD	False
UTILIZATION Asynchronous Writing Connections	Apache HTTPD	False
UTILIZATION ServerUptimeSeconds	Apache HTTPD	False
UTILIZATION Load1	Apache HTTPD	False
UTILIZATION Load5	Apache HTTPD	False
UTILIZATION ParentServerConfigGeneration	Apache HTTPD	False
UTILIZATION ParentServerMPMGeneration	Apache HTTPD	False

Apache Tomcat

Metrics are collected for the Apache Tomcat application service.

Table 8-17. Apache Tomcat

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Tomcat Server	False
Buffer Pool<InstanceName> Memory Used	Tomcat Server	False
Buffer Pool<InstanceName> Total Capacity	Tomcat Server	False
Class Loading Loaded Class Count	Tomcat Server	False
Class Loading Total Loaded Class Count	Tomcat Server	False
Class Loading Unloaded Class Count	Tomcat Server	False
File Descriptor Usage Max File Descriptor Count	Tomcat Server	False
File Descriptor Usage Open File Descriptor Count	Tomcat Server	False
Garbage Collection:<InstanceName> Total Collection Count	Tomcat Server	False
Garbage Collection:<InstanceName> Total Collection Time	Tomcat Server	True
JVM Memory Heap Memory Usage Committed Memory	Tomcat Server	False
JVM Memory Heap Memory Usage Initial Memory	Tomcat Server	False
JVM Memory Heap Memory Usage Maximum Memory	Tomcat Server	False
JVM Memory Heap Memory Usage Used Memory	Tomcat Server	False
JVM Memory Non Heap Memory Usage Committed Memory	Tomcat Server	False
JVM Memory Non Heap Memory Usage Initial Memory	Tomcat Server	False
JVM Memory Non Heap Memory Usage Maximum Memory	Tomcat Server	False
JVM Memory Non Heap Memory Usage Used Memory	Tomcat Server	False
JVM Memory Number of Object Pending Finalization Count	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Committed Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Initial Memory	Tomcat Server	False

Table 8-17. Apache Tomcat (continued)

Metric Name	Category	KPI
JVM Memory Pool:<InstanceName> Peak Usage Maximum Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Peak Usage Used Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Committed Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Initial Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Maximum Memory	Tomcat Server	False
JVM Memory Pool:<InstanceName> Usage Used Memory	Tomcat Server	False
Process CPU Usage (%)	Tomcat Server	True
System CPU Usage (%)	Tomcat Server	True
System Load Average (%)	Tomcat Server	True
Threading Thread Count	Tomcat Server	False
Uptime	Tomcat Server	True
JSP Count	Tomcat Server Web Module	False
JSP Reload Count	Tomcat Server Web Module	False
JSP Unload Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Error Count	Tomcat Server Web Module	False
Servlet:<InstanceName> Total Request Processing Time	Tomcat Server Web Module	False
Cache : Hit Count	Tomcat Server Web Module	False
Cache : Lookup Count	Tomcat Server Web Module	False
Current Thread Count	Tomcat Server Global Request Processor	True
Current Threads Busy	Tomcat Server Global Request Processor	True
errorRate	Tomcat Server Global Request Processor	False
Total Request Bytes Received	Tomcat Server Global Request Processor	False

Table 8-17. Apache Tomcat (continued)

Metric Name	Category	KPI
Total Request Bytes Sent	Tomcat Server Global Request Processor	False
Total Request Count	Tomcat Server Global Request Processor	True
Total Request Error Count	Tomcat Server Global Request Processor	True
Total Request Processing Time	Tomcat Server Global Request Processor	False

IIS Metrics

Metrics are collected for the IIS application service.

Table 8-18. IIS Metrics

Metric Name	Category	KPI
HTTP Service Request Queues<InstanceName>AppPool CurrentQueueSize	IIS HTTP Service Request Queues	True
HTTP Service Request Queues<InstanceName>AppPool RejectedRequests	IIS HTTP Service Request Queues	False
Web Services<InstanceName> Web Site Bytes Received	IIS Web Services	False
Web Services<InstanceName> Web Site Bytes Sent/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Bytes Total/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Connection Attempts/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Current Connections	IIS Web Services	False
Web Services<InstanceName> Web Site Get Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Locked Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Not Found Errors/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Post Requests/sec	IIS Web Services	False
Web Services<InstanceName> Web Site Service Uptime	IIS Web Services	False

Table 8-18. IIS Metrics (continued)

Metric Name	Category	KPI
Web Services<InstanceName> Web Site Total Bytes Sent	IIS Web Services	False
Web Services<InstanceName> Web Site Total Get Requests	IIS Web Services	True
Web Services<InstanceName> Web Site Total Post Requests	IIS Web Services	True
Web Services<InstanceName> Web Site Total Put Requests	IIS Web Services	False
Current File Cache Memory Usage (bytes)	IIS Web Services Cache	False
File Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Hits Percent (%)	IIS Web Services Cache	False
Kernel URI Cache Misses	IIS Web Services Cache	False
Total Flushed URIs	IIS Web Services Cache	False
URI Cache Hits	IIS Web Services Cache	False
URI Cache Hits Percent (%)	IIS Web Services Cache	False
URI Cache Misses	IIS Web Services Cache	False
ASP.NET<InstanceName> Application Restarts	IIS ASP.NET	True
ASP.NET<InstanceName> Request Wait Time	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Current	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Queued	IIS ASP.NET	True
ASP.NET<InstanceName> Requests Rejected	IIS ASP.NET	True
MS.NET<InstanceName> Allocated Bytes/sec	MS.NET	True
MS.NET<InstanceName> Current Queue Length	MS.NET	False
MS.NET<InstanceName> Finalization Survivors	MS.NET	False
MS.NET<InstanceName> Gen 0 Collections	MS.NET	False
MS.NET<InstanceName> Gen 0 heap size	MS.NET	False

Table 8-18. IIS Metrics (continued)

Metric Name	Category	KPI
MS.NET<InstanceName> Gen 1 Collections	MS.NET	False
MS.NET<InstanceName> Gen 1 heap size	MS.NET	False
MS.NET<InstanceName> Gen 2 Collections	MS.NET	False
MS.NET<InstanceName> Gen 2 heap size	MS.NET	False
MS.NET<InstanceName> IL Bytes Jitted / sec	MS.NET	False
MS.NET<InstanceName> Induced GC	MS.NET	False
MS.NET<InstanceName> Large Object Heap size	MS.NET	False
MS.NET<InstanceName> No of current logical Threads	MS.NET	True
MS.NET<InstanceName> No of current physical Threads	MS.NET	True
MS.NET<InstanceName> No of current recognized threads	MS.NET	False
MS.NET<InstanceName> No of Exceps Thrown / sec	MS.NET	True
MS.NET<InstanceName> No of total recognized threads	MS.NET	False
MS.NET<InstanceName> Percent Time in Jit	MS.NET	False
MS.NET<InstanceName> Pinned Objects	MS.NET	False
MS.NET<InstanceName> Stack Walk Depth	MS.NET	False
MS.NET<InstanceName> Time in RT checks	MS.NET	False
MS.NET<InstanceName> Time Loading	MS.NET	True
MS.NET<InstanceName> Total No of Contentions	MS.NET	False
MS.NET<InstanceName> Total Runtime Checks	MS.NET	True

Java Application Metrics

Metrics are collected for the Java application service.

Table 8-19. Java Application Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Java Application	False
Buffer Pool<InstanceName> Memory Used	Java Application	False
Buffer Pool<InstanceName> Total Capacity	Java Application	False
Class Loading Loaded Class Count	Java Application	True
Class Loading Total Loaded Class Count	Java Application	False
Class Loading Unloaded Class Count	Java Application	False
Garbage Collection<InstanceName> Total Collection Count	Java Application	False
Garbage Collection<InstanceName> Total Collection Time	Java Application	False
JVM Memory Heap Memory Usage Committed Memory	Java Application	False
JVM Memory Heap Memory Usage Initial Memory	Java Application	False
JVM Memory Heap Memory Usage Maximum Memory	Java Application	False
JVM Memory Heap Memory Usage Used Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Committed Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Initial Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Peak Usage Used Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Committed Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Initial Memory	Java Application	False

Table 8-19. Java Application Metrics (continued)

Metric Name	Category	KPI
JVM Memory JVM Memory Pool<InstanceName> Usage Maximum Memory	Java Application	False
JVM Memory JVM Memory Pool<InstanceName> Usage Used Memory	Java Application	False
JVM Memory Non Heap Memory Usage Committed Memory	Java Application	False
JVM Memory Non Heap Memory Usage Initial Memory	Java Application	False
JVM Memory Non Heap Memory Usage Maximum Memory	Java Application	False
JVM Memory Non Heap Memory Usage Used Memory	Java Application	False
JVM Memory Object Pending Finalization Count	Java Application	False
Uptime	Java Application	True
Threading Thread Count	Java Application	True
Process CPU Usage %	Java Application	False
System CPU Usage %	Java Application	False
System Load Average %	Java Application	False

JBoss EAP Metrics

Metrics are collected for the JBoss EAP application service.

Table 8-20. JBoss EAP Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Jboss Server	False
Buffer Pool<InstanceName> Memory Used	Jboss Server	False
Buffer Pool<InstanceName> Total Capacity	Jboss Server	False
Class Loading Loaded Class Count	Jboss Server	False
Class Loading Total Loaded Class Count	Jboss Server	False
Class Loading Unloaded Class Count	Jboss Server	False

Table 8-20. JBoss EAP Metrics (continued)

Metric Name	Category	KPI
File Descriptor Usage Max File Descriptor Count	Jboss Server	False
File Descriptor Usage Open File Descriptor Count	Jboss Server	False
Http Listener<InstanceName> Bytes Received	Jboss Server	False
Http Listener<InstanceName> Bytes Sent	Jboss Server	False
Http Listener<InstanceName> Error Count	Jboss Server	False
Http Listener<InstanceName> Request Count	Jboss Server	False
Https Listener<InstanceName> Bytes Received	Jboss Server	False
Https Listener<InstanceName> Bytes Sent	Jboss Server	False
Https Listener<InstanceName> Error Count	Jboss Server	False
Https Listener<InstanceName> Request Count	Jboss Server	False
Process CPU Usage (%)	Jboss Server	False
System CPU Usage (%)	Jboss Server	False
System Load Average (%)	Jboss Server	False
Threading Daemon Thread Count	Jboss Server	False
Threading Peak Thread Count	Jboss Server	False
Threading Thread Count	Jboss Server	False
Threading Total Started Thread Count	Jboss Server	False
Uptime	Jboss Server	False
UTILIZATION Heap Memory Usage	Jboss Server	False
Garbage Collection<InstanceName> Total Collection Count	Jboss JVM Garbage Collector	False
Garbage Collection<InstanceName> Total Collection Time	Jboss JVM Garbage Collector	False
JVM Memory Heap Memory Usage Committed Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Initial Memory	Jboss JVM Memory	False

Table 8-20. JBoss EAP Metrics (continued)

Metric Name	Category	KPI
JVM Memory Heap Memory Usage Maximum Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Used Memory	Jboss JVM Memory	True
JVM Memory Non Heap Memory Usage Committed Memory	Jboss JVM Memory	False
JVM Memory Non Heap Memory Usage Initial Memory	Jboss JVM Memory	False
JVM Memory Non Heap Memory Usage Maximum Memory	Jboss JVM Memory	False
JVM Memory Non Heap Memory Usage Used Memory	Jboss JVM Memory	False
JVM Memory Object Pending Finalization Count	Jboss JVM Memory	True
UTILIZATION Active Count	Jboss Datasource Pool	False
UTILIZATION Available Count	Jboss Datasource Pool	False
JVM Memory Pool<InstanceName> Collection Usage Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Collection Usage Initial Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Collection Usage Used Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Collection Usage Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Peak Usage Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Peak Usage Initial Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Peak Usage Used Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Usage Committed Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Usage Initial Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Usage Maximum Memory	Jboss JVM Memory Pool	False
JVM Memory Pool<InstanceName> Usage Used Memory	Jboss JVM Memory Pool	False

Hyper-V Metrics

Metrics are collected for the Hyper-V application service.

Table 8-21. Hyper-V Metrics

Metric Name	Category	KPI
VM:Hyper-V Virtual Machine Health Summary Health Critical	HyperV	False
VM<instanceName> Physical Memory	HyperV	False
VM<instanceName> Hv VP 0 Total Run Time	HyperV	False
VM<instanceName> Bytes Received	HyperV	False
VM<instanceName> Bytes Sent	HyperV	False
VM<instanceName> Error Count	HyperV	False
VM<instanceName> Latency	HyperV	False
VM<instanceName> Queue Length	HyperV	False
VM<instanceName> Throughput	HyperV	False
CPU<instanceName> Idle Time	HyperV	True
CPU<instanceName> Processor Time	HyperV	True
CPU<instanceName> User Time	HyperV	True
Disk<instanceName> Avg Disk Queue Length	HyperV	False
Disk<instanceName> Idle Time	HyperV	False
Disk<instanceName> Read Time	HyperV	True
Disk<instanceName> Write Time	HyperV	True
Process<instanceName> Private Bytes	HyperV	False
Process<instanceName> Processor Time	HyperV	False
Process<instanceName> Thread Count	HyperV	False
Process<instanceName> User Time	HyperV	False
System Processes	HyperV	False
System Processor Queue Length	HyperV	False
System System UpTime	HyperV	False
Memory Available Bytes	HyperV	False

Table 8-21. Hyper-V Metrics (continued)

Metric Name	Category	KPI
Memory Cache Bytes	HyperV	False
Memory Cache Faults	HyperV	False
Memory Pages	HyperV	False
Network<instanceName> Packets Outbound Error	HyperV	False
Network<instanceName> Packets Received Error	HyperV	False

Oracle Database Metrics

Metrics are collected for the Oracle database application service.

Oracle database cannot be activated on Linux platforms.

Table 8-22. Oracle Database Metrics

Metric Name	Category	KPI
Utilization Active Sessions	OracleDB	True
Utilization Buffer CacheHit Ratio	OracleDB	False
Utilization Cursor CacheHit Ratio	OracleDB	False
Utilization Database Wait Time	OracleDB	False
Utilization Disk Sort persec	OracleDB	False
Utilization Enqueue Timeouts Persec	OracleDB	False
Utilization Global Cache Blocks Corrupted	OracleDB	False
Utilization Global Cache Blocks Lost	OracleDB	False
Utilization Library CacheHit Ratio	OracleDB	False
Utilization Logon persec	OracleDB	True
Utilization Memory Sorts Ratio	OracleDB	True
Utilization Rows persort	OracleDB	False
Utilization Service Response Time	OracleDB	False
Utilization Session Count	OracleDB	True
Utilization Session Limit	OracleDB	False
Utilization Shared Pool Free	OracleDB	False
Utilization Temp Space Used	OracleDB	False

Table 8-22. Oracle Database Metrics (continued)

Metric Name	Category	KPI
Utilization Total Sorts persec	OracleDB	False
Utilization Physical Read Bytes Perc	OracleDB	False
Utilization Physical Read IO Requests Perc	OracleDB	False
Utilization Physical Read Total Bytes Persec	OracleDB	False
Utilization Physical Reads Persec	OracleDB	True
Utilization Physical Reads Per Txn	OracleDB	False
Utilization Physical Write Bytes Perc	OracleDB	False
Utilization Physical Write IO Requests Perc	OracleDB	False
Utilization Physical Write Total Bytes Perc	OracleDB	False
Utilization Physical Writes Perc	OracleDB	True
Utilization Physical Writes Per Txn	OracleDB	False
Utilization User Commits Percentage	OracleDB	False
Utilization User Commits Perc	OracleDB	False
Utilization User Rollbacks Percentage	OracleDB	False
Utilization User Rollbacks persec	OracleDB	True
Utilization User Transaction Persec	OracleDB	False
Utilization Database Time Perc	OracleDB	False

Cassandra Database Metrics

Metrics are collected for the Cassandra database application service.

Table 8-23. Cassandra Database Metrics

Metric Name	Category	KPI
Cache<InstanceName> Capacity	Cassandra	False
Cache<InstanceName> Entries	Cassandra	True
Cache<InstanceName> HitRate	Cassandra	True
Cache<InstanceName> Requests	Cassandra	True
Cache<InstanceName> Size	Cassandra	False

Table 8-23. Cassandra Database Metrics (continued)

Metric Name	Category	KPI
ClientRequest<InstanceName> Failures	Cassandra	False
ClientRequest<InstanceName> Latency	Cassandra	False
ClientRequest<InstanceName> Timeouts	Cassandra	False
ClientRequest<InstanceName> Total Latency	Cassandra	False
ClientRequest<InstanceName> Unavailables	Cassandra	False
CommitLog Pending Tasks	Cassandra	False
CommitLog Total Commit Log Size	Cassandra	False
Compaction Bytes Compacted	Cassandra	False
Compaction Completed Tasks	Cassandra	False
Compaction Pending Tasks	Cassandra	False
Compaction Total Compactions Completed	Cassandra	False
Connected Native Clients	Cassandra	False
HeapMemoryUsage committed	Cassandra	False
HeapMemoryUsage init	Cassandra	False
HeapMemoryUsage max	Cassandra	False
HeapMemoryUsage used	Cassandra	False
NonHeapMemoryUsage committed	Cassandra	False
NonHeapMemoryUsage init	Cassandra	False
NonHeapMemoryUsage max	Cassandra	False
NonHeapMemoryUsage used	Cassandra	False
ObjectPendingFinalizationCount	Cassandra	False
Storage Exceptions Count	Cassandra	False
Storage Load Count	Cassandra	False
Table<InstanceName> Coordinator Read Latency	Cassandra	False
Table<InstanceName> Live Diskpace Used	Cassandra	False

Table 8-23. Cassandra Database Metrics (continued)

Metric Name	Category	KPI
Table<InstanceName> Read Latency	Cassandra	False
Table<InstanceName> Total Diskspace Used	Cassandra	False
Table<InstanceName> Total Read Latency	Cassandra	False
Table<InstanceName> Total Write Latency	Cassandra	False
Table<InstanceName> Write Latency	Cassandra	False
ThreadPools<InstanceName> Active Tasks	Cassandra	False
ThreadPools<InstanceName> Currently Blocked Tasks	Cassandra	False
ThreadPools<InstanceName> Pending Tasks	Cassandra	False

MongoDB Metrics

Metrics are collected for the MongoDB application service.

Table 8-24. MongoDB Metrics

Metric Name	Category	KPI
UTILIZATION Active Reads	MongoDB	True
UTILIZATION Active Writes	MongoDB	True
UTILIZATION Connections Available	MongoDB	False
UTILIZATION Connections Total Created	MongoDB	False
UTILIZATION Current Connections	MongoDB	True
UTILIZATION Cursor Timed Out	MongoDB	True
UTILIZATION Deletes Per Sec	MongoDB	False
UTILIZATION Document Inserted	MongoDB	False
UTILIZATION Document Deleted	MongoDB	False
UTILIZATION Flushes Per Sec	MongoDB	False
UTILIZATION Inserts Per Sec	MongoDB	False
UTILIZATION Net Input Bytes	MongoDB	False
UTILIZATION Open Connections	MongoDB	True

Table 8-24. MongoDB Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Page Faults Per Second	MongoDB	False
UTILIZATION Net Output Bytes	MongoDB	False
UTILIZATION Queries Per Sec	MongoDB	False
UTILIZATION Queued Reads	MongoDB	True
UTILIZATION Queued Writes	MongoDB	True
UTILIZATION Total Available	MongoDB	False
UTILIZATION Total Deletes Per Sec	MongoDB	False
UTILIZATION Total Passes Per Sec	MongoDB	False
UTILIZATION Total Refreshing	MongoDB	False
UTILIZATION Updates Per Sec	MongoDB	False
UTILIZATION Volume Size MB	MongoDB	False
UTILIZATION Collection Stats	MongoDB DataBases	False
UTILIZATION Data Index Stats	MongoDB DataBases	True
UTILIZATION Data Indexes	MongoDB DataBases	False
UTILIZATION Data Size Stats	MongoDB DataBases	True
UTILIZATION Average Object Size stats	MongoDB DataBases	False
UTILIZATION Num Extents Stats	MongoDB DataBases	False

MS Exchange Server Metrics

Metrics are collected for the MS Exchange Server application service.

Table 8-25. MS Exchange Server Metrics

Metric Name	Category	KPI
Active Manager Server Active Manager Role	MS Exchange	False
Active Manager Server Database State Info Writes per second	MS Exchange	False
Active Manager Server GetServerForDatabase Server-Side Calls	MS Exchange	False
Active Manager Server Server-Side Calls per second	MS Exchange	True

Table 8-25. MS Exchange Server Metrics (continued)

Metric Name	Category	KPI
Active Manager Server Total Number of Databases	MS Exchange	True
ActiveSync Average Request Time	MS Exchange	True
ActiveSync Current Requests	MS Exchange	False
ActiveSync Mailbox Search Total	MS Exchange	False
ActiveSync Ping Commands Pending	MS Exchange	False
ActiveSync Requests per second	MS Exchange	True
ActiveSync Sync Commands per second	MS Exchange	True
ASP.NET Application Restarts	MS Exchange	False
ASP.NET Request Wait Time	MS Exchange	True
ASP.NET Worker Process Restarts	MS Exchange	False
Autodiscover Service Requests per second	MS Exchange	True
Availability Service Average Time to Process a Free Busy Request	MS Exchange	True
Outlook Web Access Average Search Time	MS Exchange	True
Outlook Web Access Requests per second	MS Exchange	False
Outlook Web Access Current Unique Users	MS Exchange	False
Performance Database Cache Hit (%)	MS Exchange Database	False
Performance Database Page Fault Stalls per second	MS Exchange Database	True
Performance I/O Database Reads Average Latency	MS Exchange Database	True
Performance I/O Database Writes Average Latency	MS Exchange Database	True
Performance I/O Log Reads Average Latency	MS Exchange Database	False
Performance I/O Log Writes Average Latency	MS Exchange Database	False
Performance Log Record Stalls per second	MS Exchange Database	False
Performance Log Threads Waiting	MS Exchange Database	False

Table 8-25. MS Exchange Server Metrics (continued)

Metric Name	Category	KPI
Performance\I/O Database Reads Average Latency	MS Exchange Database Instance	False
Performance\I/O Database Writes Average Latency	MS Exchange Database Instance	False
Performance\Log Record Stalls per second	MS Exchange Database Instance	False
Performance\Log Threads Waiting	MS Exchange Database Instance	False
Performance\LDAP Read Time	MS Exchange Domain Controller	False
Performance\LDAP Search Time	MS Exchange Domain Controller	False
Performance\LDAP Searches Timed Out per minute	MS Exchange Domain Controller	False
Performance\Long Running LDAP Operations per minute	MS Exchange Domain Controller	False
Performance\Connection Attempts per second	MS Exchange Web Server	True
Performance\Current Connections	MS Exchange Web Server	False
Performance\Other Request Methods per second	MS Exchange Web Server	False
Process\Handle Count	MS Exchange Windows Service	False
Process\Memory Allocated	MS Exchange Windows Service	False
Process\Processor Time (%)	MS Exchange Windows Service	True
Process\Thread Count	MS Exchange Windows Service	False
Process\Virtual Memory Used	MS Exchange Windows Service	False
Process\Working Set	MS Exchange Windows Service	False

MS SQL Metrics

Metrics are collected for the MS SQL application service.

Table 8-26. MS SQL Metrics

Metric Name	Category	KPI
CPU<InstanceName>\CPU Usage (%)	Microsoft SQL Server	False
Database IO\Rows Reads Bytes/Sec	Microsoft SQL Server	False
Database IO\Rows Reads/Sec	Microsoft SQL Server	False
Database IO\Rows Writes Bytes/Sec	Microsoft SQL Server	False

Table 8-26. MS SQL Metrics (continued)

Metric Name	Category	KPI
Database IO Rows Writes/Sec	Microsoft SQL Server	False
Performance Access Methods Full Scans per second	Microsoft SQL Server	False
Performance Access Methods Index Searches	Microsoft SQL Server	False
Performance Access Methods Page Splits per second	Microsoft SQL Server	False
Performance Broker Activation Stored Procedures Invoked per second	Microsoft SQL Server	False
Performance Buffer Manager Buffer cache hit ratio (%)	Microsoft SQL Server	True
Performance Buffer Manager Checkpoint Pages/sec	Microsoft SQL Server	True
Performance Buffer Manager Lazy writes per second	Microsoft SQL Server	True
Performance Buffer Manager Page life expectancy	Microsoft SQL Server	True
Performance Buffer Manager Page lookups per second	Microsoft SQL Server	False
Performance Buffer Manager Page reads per second	Microsoft SQL Server	False
Performance Buffer Manager Page writes per second	Microsoft SQL Server	False
Performance Databases Active Transactions	Microsoft SQL Server	True
Performance Databases Data File(s) Size	Microsoft SQL Server	True
Performance Databases Log Bytes Flushed/Sec	Microsoft SQL Server	False
Performance Databases Log File(s) Size	Microsoft SQL Server	False
Performance Databases Log File(s) Used Size	Microsoft SQL Server	False
Performance Databases Log Flush Wait Time	Microsoft SQL Server	False
Performance Databases Log Flushes per second	Microsoft SQL Server	False
Performance Databases Transactions per second	Microsoft SQL Server	False

Table 8-26. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance\Databases\Write Transactions per second	Microsoft SQL Server	False
Performance\Databases\XTP Memory Used	Microsoft SQL Server	False
Performance\General Statistics\Active temp Tables	Microsoft SQL Server	False
Performance\General Statistics\Logins per second	Microsoft SQL Server	False
Performance\General Statistics\Logouts per second	Microsoft SQL Server	False
Performance\General Statistics\Processes Blocked	Microsoft SQL Server	False
Performance\General Statistics\Temp Tables Creation Rate	Microsoft SQL Server	False
Performance\General Statistics\User Connections	Microsoft SQL Server	False
Performance\Locks\Average Wait Time	Microsoft SQL Server	False
Performance\Locks\Lock Requests per second	Microsoft SQL Server	False
Performance\Locks\Lock Wait Time	Microsoft SQL Server	True
Performance\Locks\Lock Waits per second	Microsoft SQL Server	True
Performance\Locks\Number of Deadlocks per second	Microsoft SQL Server	True
Performance\Memory Manager\Connection Memory	Microsoft SQL Server	False
Performance\Memory Manager\Lock Memory	Microsoft SQL Server	False
Performance\Memory Manager\Log Pool Memory	Microsoft SQL Server	False
Performance\Memory Manager\Memory Grants Pending	Microsoft SQL Server	True
Performance\Memory Manager\SQL Cache Memory	Microsoft SQL Server	False
Performance\Memory Manager\Target Server Memory	Microsoft SQL Server	True
Performance\Memory Manager\Total Server Memory	Microsoft SQL Server	True
Performance\Resource Pool Stats\internal\Active memory grant amount	Microsoft SQL Server	False

Table 8-26. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance Resource Pool Stats internal CPU Usage Percentage (%)	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read Bytes per second	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read IO	Microsoft SQL Server	False
Wait Stats:<InstanceName> Wait Time (ms)	Microsoft SQL Server	False
Wait Stats<InstanceName> Number of Waiting tasks (ms)	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Read IO Throttled Per Second	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Write Bytes per second (Bps)	Microsoft SQL Server	False
Performance Resource Pool Stats internal Disk Write IO Throttled per second	Microsoft SQL Server	False
Performance Resource Pool Stats internal Used Memory	Microsoft SQL Server	False
Performance SQL Statistics Batch Requests Per Second	Microsoft SQL Server	False
Performance SQL Statistics SQL Compilations per second	Microsoft SQL Server	False
Performance SQL Statistics SQL Re-Compilations per second	Microsoft SQL Server	False
Performance Transactions Free space in tempdb (KB)	Microsoft SQL Server	False
Performance Transactions Transactions	Microsoft SQL Server	False
Performance Transactions Version Store Size (KB)	Microsoft SQL Server	False
Performance User Settable Counter User Counter 0 to 10	Microsoft SQL Server	False
Performance Workload Group Stats internal Active Requests	Microsoft SQL Server	False
Performance Workload Group Stats internal Blocked Tasks	Microsoft SQL Server	False
Performance Workload Group Stats internal CpU Usage (%)	Microsoft SQL Server	False

Table 8-26. MS SQL Metrics (continued)

Metric Name	Category	KPI
Performance\Workload Group Stats\internal\Queued Requests	Microsoft SQL Server	False
Performance\Workload Group Stats\internal\Request Completed/sec	Microsoft SQL Server	False

There are no metrics collected for Microsoft SQL Server Database.

MySQL Metrics

Metrics are collected for the MySQL application service.

Table 8-27. MySQL Metrics

Metric Name	Category	KPI
Aborted connection count	MySQL	True
Connection count	MySQL	True
Event wait average time	MySQL	False
Event wait count	MySQL	False
Binary Files\Binary Files Count	MySQL	False
Binary Files\Binary Size Bytes	MySQL	False
Global Status\Aborted Clients	MySQL	False
Global Status\Binlog Cache Disk Use	MySQL	False
Global Status\Bytes Received	MySQL	False
Global Status\Bytes Sent	MySQL	False
Global Status\Connection Errors Accept	MySQL	False
Global Status\Connection Errors Internal	MySQL	False
Global Status\Connection Errors Max Connections	MySQL	False
Global Status\Queries	MySQL	False
Global Status\Threads Cached	MySQL	False
Global Status\Threads Connected	MySQL	False
Global Status\Threads Running	MySQL	False
Global Status\Uptime	MySQL	False
Global Variables\Delayed Insert Limit	MySQL	False

Table 8-27. MySQL Metrics (continued)

Metric Name	Category	KPI
Global Variables Delayed Insert Timeout	MySQL	False
Global Variables Delayed Queue Size	MySQL	False
Global Variables Max Connect Errors	MySQL	False
Global Variables Max Connections	MySQL	False
Global Variables Max Delayed Threads	MySQL	False
Global Variables Max Error Count	MySQL	False
InnoDB All deadlock count	MySQL	False
InnoDB Buffer Pool Bytes Data	MySQL	False
InnoDB Buffer Pool Bytes Data	MySQL	False
InnoDB Buffer Pool Bytes Dirty	MySQL	False
InnoDB Buffer Pool Dump Status	MySQL	False
InnoDB Buffer Pool Load Status	MySQL	False
InnoDB Buffer Pool Pages Data	MySQL	False
InnoDB Buffer Pool Pages Dirty	MySQL	False
InnoDB Buffer Pool Pages Flushed	MySQL	False
InnoDB Buffer pool size	MySQL	True
InnoDB Checksums	MySQL	False
InnoDB Open file count	MySQL	False
InnoDB Row lock average time	MySQL	False
InnoDB Row lock current waits	MySQL	False
InnoDB Row lock maximum time	MySQL	False
InnoDB Row lock time	MySQL	False
InnoDB Row lock waits	MySQL	True
InnoDB Table lock count	MySQL	False
Performance Table IO Waits IO Waits Total Delete	MySQL	False
Performance Table IO Waits IO Waits Total Fetch	MySQL	False

Table 8-27. MySQL Metrics (continued)

Metric Name	Category	KPI
Performance Table IO Waits IO Waits Total Insert	MySQL	False
Performance Table IO Waits IO Waits Total Update	MySQL	False
Process List Connections	MySQL	False
IO waits average time	MySQL Database	False
IO waits count	MySQL Database	True
Read high priority average time	MySQL Database	False
Read high priority count	MySQL Database	False
Write concurrent insert average time	MySQL Database	False
Write concurrent insert count	MySQL Database	False

NGINX Metrics

Metrics are collected for the NGINX application service.

Table 8-28. NGINX Metrics

Metric Name	Category	KPI
HTTP Status Info Accepts	Nginx	True
HTTP Status Info Active connections	Nginx	False
HTTP Status Info Handled	Nginx	True
HTTP Status Info Reading	Nginx	False
HTTP Status Info Requests	Nginx	False
HTTP Status Info Waiting	Nginx	True
HTTP Status Info Writing	Nginx	False

NTPD Metrics

Metrics are collected for the NTPD application service.

Table 8-29. NTPD Metrics

Metric Name	Category	KPI
ntpd delay	Network Time Protocol	True
ntpd jitter	Network Time Protocol	True

Table 8-29. NTPD Metrics (continued)

Metric Name	Category	KPI
ntpd offset	Network Time Protocol	True
ntpd poll	Network Time Protocol	False
ntpd reach	Network Time Protocol	True
ntpd when	Network Time Protocol	False

Oracle Weblogic Metrics

Metrics are collected for the Oracle Weblogic application service.

Table 8-30. Oracle Weblogic Metrics

Metric Name	Category	KPI
UTILIZATION Process Cpu Load	Oracle WebLogic Server	True
UTILIZATION System Cpu Load	Oracle WebLogic Server	False
UTILIZATION System Load Average	Oracle WebLogic Server	False
UTILIZATION Collection Time	Weblogic Garbage Collector	True
UTILIZATION Connections HighCount	Weblogic JMS Runtime	True
UTILIZATION JMS Servers TotalCount	Weblogic JMS Runtime	False
UTILIZATION Active Total Count Used	Weblogic JTA Runtime	False
UTILIZATION Active Transactions TotalCount	Weblogic JTA Runtime	False
UTILIZATION Transaction Abandoned TotalCount	Weblogic JTA Runtime	True
UTILIZATION Transaction RolledBack App TotalCount	Weblogic JTA Runtime	True
UTILIZATION Heap Memory Usage	Weblogic JVM Memory	True
UTILIZATION Non Heap Memory Usage	Weblogic JVM Memory	False
UTILIZATION Peak Usage	Weblogic JVM Memory Pool	True
UTILIZATION Usage	Weblogic JVM Memory Pool	False
UTILIZATION UpTime	Weblogic JVM Runtime	False

Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

Table 8-31. Pivotal TC Server Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Pivotal TC Server	False
Buffer Pool<InstanceName> Memory Used	Pivotal TC Server	False
Buffer Pool<InstanceName> Total Capacity	Pivotal TC Server	False
Class Loading Loaded Class Count	Pivotal TC Server	False
Class Loading Total Loaded Class Count	Pivotal TC Server	False
Class Loading Unloaded Class Count	Pivotal TC Server	False
File Descriptor Usage Max File Descriptor Count	Pivotal TC Server	False
File Descriptor Usage Open File Descriptor Count	Pivotal TC Server	False
Garbage Collection:<InstanceName> Total Collection Count	Pivotal TC Server	False
Garbage Collection:<InstanceName> Total Collection Time	Pivotal TC Server	False
Process CPU Usage (%)	Pivotal TC Server	True
JVM Memory Heap Memory Usage Committed Memory	Pivotal TC Server	True
JVM Memory Heap Memory Usage Initial Memory	Pivotal TC Server	False
JVM Memory Heap Memory Usage Maximum Memory	Pivotal TC Server	False
JVM Memory Heap Memory Usage Used Memory	Pivotal TC Server	True
JVM Memory Non Heap Memory Usage Committed Memory	Pivotal TC Server	True
JVM Memory Non Heap Memory Usage Initial Memory	Pivotal TC Server	False
JVM Memory Non Heap Memory Usage Maximum Memory	Pivotal TC Server	False
JVM Memory Non Heap Memory Usage Used Memory	Pivotal TC Server	True
JVM Memory Number of Object Pending Finalization Count	Pivotal TC Server	True
JVM Memory Pool:<InstanceName> Peak Usage Committed Memory	Pivotal TC Server	False

Table 8-31. Pivotal TC Server Metrics (continued)

Metric Name	Category	KPI
JVM Memory Pool:<InstanceName> Peak Usage Initial Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Peak Usage Maximum Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Peak Usage Used Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Usage Committed Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Usage Initial Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Usage Maximum Memory	Pivotal TC Server	False
JVM Memory Pool:<InstanceName> Usage Used Memory	Pivotal TC Server	False
Process CPU Usage (%)	Pivotal TC Server	True
System CPU Usage (%)	Pivotal TC Server	True
Uptime	Pivotal TC Server	True
Threading Thread Count	Pivotal TC Server	False
System Load Average	Pivotal TC Server	False
Current Thread Count	Pivotal TC Server Thread Pool	False
Current Threads Busy	Pivotal TC Server Thread Pool	True
Total Request Bytes Received	Pivotal TC Server Thread Pool	False
Total Request Bytes Sent	Pivotal TC Server Thread Pool	False
Total Request Count	Pivotal TC Server Thread Pool	True
Total Request Error Count	Pivotal TC Server Thread Pool	True
Total Request Processing Time	Pivotal TC Server Thread Pool	True
JSP Count	Pivotal TC Server Web Module	False
JSP Reload Count	Pivotal TC Server Web Module	False
JSP Unload Count	Pivotal TC Server Web Module	False

PostgreSQL

Metrics are collected for the PostgreSQL application service.

Table 8-32. PostgreSQL

Metric Name	Category	KPI
Buffers Buffers Allocated	PostgreSQL	False
Buffers Buffers Written by Backend	PostgreSQL	True
Buffers Buffers Written by Background Writer	PostgreSQL	True
Buffers Buffers Written During Checkpoints	PostgreSQL	True
Buffers fsync Call Executed by Backend	PostgreSQL	False
Checkpoints Checkpoints sync time	PostgreSQL	False
Checkpoints Checkpoints write time	PostgreSQL	False
Checkpoints Requested checkpoints performed count	PostgreSQL	False
Checkpoints Scheduled checkpoints performed count	PostgreSQL	False
Clean scan stopped count	PostgreSQL	False
Disk Blocks Blocks Cache Hits	PostgreSQL Database	False
Disk Blocks Blocks Read	PostgreSQL Database	False
Disk Blocks Blocks Read Time	PostgreSQL Database	False
Disk Blocks Blocks Write Time	PostgreSQL Database	False
Statistics Backends Connected	PostgreSQL Database	False
Statistics Data Written by Queries	PostgreSQL Database	True
Statistics Deadlocks Detected	PostgreSQL Database	True
Statistics Queries Cancelled	PostgreSQL Database	True
Statistics Temp Files Created by Queries	PostgreSQL Database	False
Transactions Transactions Committed	PostgreSQL Database	True
Transactions Transactions Rolled Back	PostgreSQL Database	True
Tuples Tuples Deleted	PostgreSQL Database	True
Tuples Tuples Fetched	PostgreSQL Database	True
Tuples Tuples Inserted	PostgreSQL Database	True
Tuples Tuples Returned	PostgreSQL Database	True
Tuples Tuples Updated	PostgreSQL Database	True

RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

Table 8-33. RabbitMQ Metrics

Metric Name	Category	KPI
CPU Limit	RabbitMQ	False
CPU Used	RabbitMQ	True
Disk Free	RabbitMQ	False
Disk Free limit	RabbitMQ	False
FileDescriptor Total	RabbitMQ	False
FileDescriptor Used	RabbitMQ	False
Memory Limit	RabbitMQ	False
Memory Used	RabbitMQ	True
Messages Aked	RabbitMQ	False
Messages Delivered	RabbitMQ	False
Messages Delivered get	RabbitMQ	False
Messages Published	RabbitMQ	False
Messages Ready	RabbitMQ	False
Messages Unacked	RabbitMQ	False
Socket Limit	RabbitMQ	False
Socket Used	RabbitMQ	True
UTILIZATION Channels	RabbitMQ	True
UTILIZATION Connections	RabbitMQ	True
UTILIZATION Consumers	RabbitMQ	True
UTILIZATION Exchanges	RabbitMQ	True
UTILIZATION Messages	RabbitMQ	True
UTILIZATION Queues	RabbitMQ	True
Messages Publish in	RabbitMQ Exchange	False
Messages Publish out	RabbitMQ Exchange	False
Consumer Utilisation	RabbitMQ Queue	False
Consumers	RabbitMQ Queue	False

Table 8-33. RabbitMQ Metrics (continued)

Metric Name	Category	KPI
Memory	RabbitMQ Queue	False
Messages Ack	RabbitMQ Queue	False
Messages Ack rate	RabbitMQ Queue	False
Messages Deliver	RabbitMQ Queue	False
Messages Deliver get	RabbitMQ Queue	False
Messages Persist	RabbitMQ Queue	False
Messages Publish	RabbitMQ Queue	False
Messages Publish rate	RabbitMQ Queue	False
Messages Ram	RabbitMQ Queue	False
Messages Ready	RabbitMQ Queue	False
Messages Redeliver	RabbitMQ Queue	False
Messages Redeliver rate	RabbitMQ Queue	False
Messages Space	RabbitMQ Queue	False
Messages Unack	RabbitMQ Queue	False
Messages Unacked	RabbitMQ Queue	False
Messages	RabbitMQ Queue	False

There are no metrics collected for RabbitMQ Virtual Host.

Riak Metrics

Metrics are collected for the Riak application service.

Table 8-34. Riak Metrics

Metric Name	Category	KPI
UTILIZATION CPU Average	Riak KV	False
UTILIZATION Memory Processes	Riak KV	False
UTILIZATION Memory Total	Riak KV	False
UTILIZATION Node GETs	Riak KV	True
UTILIZATION Node GETs Total	Riak KV	False
UTILIZATION Node PUTs	Riak KV	True
UTILIZATION Node PUTs Total	Riak KV	False

Table 8-34. Riak Metrics (continued)

Metric Name	Category	KPI
UTILIZATION PBC Active	Riak KV	True
UTILIZATION PBC Connects	Riak KV	True
UTILIZATION Read Repairs	Riak KV	True
UTILIZATION vNODE Index Reads	Riak KV	True
UTILIZATION vNODE Index Writes	Riak KV	True

Sharepoint Metrics

Metrics are collected for the Sharepoint application service.

Table 8-35. Sharepoint Metrics

Metric Name	Category	KPI
Sharepoint Foundation Active Threads	SharePoint Server	True
Sharepoint Foundation Current Page Requests	SharePoint Server	False
Sharepoint Foundation Executing SQL Queries	SharePoint Server	False
Sharepoint Foundation Executing Time/Page Request	SharePoint Server	True
Sharepoint Foundation Incoming Page Requests Rate	SharePoint Server	False
Sharepoint Foundation Object Cache Hit Count	SharePoint Server	False
Sharepoint Foundation Reject Page Requests Rate	SharePoint Server	False
Sharepoint Foundation Responded Page Requests Rate	SharePoint Server	True
SQL query executing time	SharePoint Server	False
Network Received Data Rate	SharePoint Web Server	True
Network Sent Data Rate	SharePoint Web Server	True
Process Processor Time (%)	SharePoint Windows Service	False
Process Threads	SharePoint Windows Service	False

WebSphere Metrics

Metrics are collected for the WebSphere application service.

Table 8-36. WebSphere Metrics

Metric Name	Category	KPI
Thread Pool Active Count Current	Thread Pool	False
Thread Pool Active Count High	Thread Pool	False
Thread Pool Active Count Low	Thread Pool	False
Thread Pool Active Count Lower	Thread Pool	False
Thread Pool Active Count Upper	Thread Pool	False
JDBC Close Count	JDBC	False
JDBC Create Count	JDBC	False
JDBC JDBC Pool Size Average	JDBC	False
JDBC JDBC Pool Size Current	JDBC	False
JDBC JDBC Pool Size Lower	JDBC	False
JDBC JDBC Pool Size Upper	JDBC	False
Garbage Collection<InstanceName> Total Collection Count	WebSphere	False
Garbage Collection<InstanceName> Total Collection Time	WebSphere	False
JVM Memory Heap Memory Usage Committed Memory	WebSphere	False
JVM Memory Heap Memory Usage Initial Memory	WebSphere	False
JVM Memory Heap Memory Usage Maximum Memory	WebSphere	False
JVM Memory Heap Memory Usage Used Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Committed Memory	WebSphere	False

Table 8-36. WebSphere Metrics (continued)

Metric Name	Category	KPI
JVM Memory Non Heap Memory Usage Initial Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Maximum Memory	WebSphere	False
JVM Memory Non Heap Memory Usage Used Memory	WebSphere	False
JVM Memory Number of Object Pending Finalization Count	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Committed Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Initial Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Maximum Memory	WebSphere	False
JVM Memory Pool<InstanceName> Peak Usage Used Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Committed Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Initial Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Maximum Memory	WebSphere	False
JVM Memory Pool<InstanceName> Usage Used Memory	WebSphere	False
Process Cpu Load	WebSphere	False
System Cpu Load	WebSphere	False
System Load Average	WebSphere	False

VeloCloud Application Service Metrics

Metrics are collected for application services supported by VeloCloud.

VeloCloud Gateway Metrics

Metrics are collected for the VeloCloud Gateway.

Table 8-37. VeloCloud Gateway Metrics

Component	Metrics
DPDK	DPDK:mbuf pool free
NAT	NAT Active Flows (%)
	NAT Active Flows
	NAT Active Routes
	NAT Active Routes Used (%)
	NAT Connected Peers
	NAT NAT Entries
NTP Server	NTP Server:ntp.ubuntu.com offset value
Summary	Summary Active Tunnels Count (%)
	Summary Average Packets Dropped
	Summary Average wMarkDrop
	Summary BGP Enabled VRFs
	Summary BGP Neighbors
	Summary CLR Count
	Summary Connected Edges
	Summary NAT
	Summary SSH Failed Login
	Summary Unstable Path Percentage
	Summary VMCP CTRL Drop Count
	Summary VMCP TX Drop Count
VC Queue	VC Queue ipv4_bh packet drop
VCMP Tunnel	VCMP Tunnel ctrl_0 packet drop
	VCMP Tunnel ctrl_1 packet drop

Table 8-37. VeloCloud Gateway Metrics (continued)

Component	Metrics
	VCMP Tunnel data_0 packet drop
	VCMP Tunnel data_1 packet drop
	VCMP Tunnel init packet drop

VeloCloud Orchestrator Metrics

Metrics are collected for the VeloCloud Orchestrator.

Table 8-38. VeloCloud Orchestrator Metrics

Component	Metrics
General	General Free Memory (%)
	General Status

Metrics - Nginx

Metrics are collected for the VeloCloud Nginx.

Table 8-39. Nginx Metrics

Component	Metrics
HTTP Status Info	HTTP Status Info Accepts
	HTTP Status Info Active Connections
	HTTP Status Info Handled
	HTTP Status Info Reading
	HTTP Status Info Requests
	HTTP Status Info Waiting
	HTTP Status Info Writing

Metrics - Redis

Metrics are collected for the VeloCloud Redis.

Table 8-40. Redis Metrics

Component	Metrics
Publish Subscribe.	Publish Subscribe Channels
Total	Total Commands Processed
	Total Connections Received

Table 8-40. Redis Metrics (continued)

Component	Metrics
Used	Used CPU
	Used Memory
	Used Peak Memory

Metrics - ClickHouse

Metrics are collected for the VeloCloud Clickhouse.

Table 8-41. Clickhouse Metrics

Component	Metrics
Background	Background Pool Task
Buffer	Buffers Allocation (Bytes)
	Buffers Compressed Read Buffer (Bytes)
	Buffers Compressed Read Buffer Blocks
	Buffers IO Allocation (Bytes)
	Buffers Storage Buffer (Bytes)
	Buffers Storage Buffer Rows
Events	Events Context Lock
	Events Disk Write Elapsed (µs)
	Events File Open
	Events Function Execute
	Events Hard Page Faults
	Events Lock Readers Wait (µs)
	Events OS IO wait (ms)
	Events OS Write (Bytes)
	Events Query
	Events Readers Wait (ms)
	Events Real Time
	Events Soft Page Faults (µs)
	Events System Time (µs)
	Events User Time (µs)

Table 8-41. Clickhouse Metrics (continued)

Component	Metrics
Global Thread	Global Global Thread
	Global Global Thread Active
Local Thread	Local Local Thread
	Local Local Thread Active
Replicas	Replicas Max Absolute Delay
	Replicas Max Insert In Queue
	Replicas Max Merge In Queue
	Replicas Max Queue Size
	Replicas Max Relative Delay
	Replicas Total Insert In Queue
	Replicas Total Merge Queues
	Replicas Total Queue Size
Summary	Summary Background Pool Task
	Summary Dict Cache Requests
	Summary File Open Writes
	Summary Merge
	Summary Number of Databases
	Summary Number of Distributed Send
	Summary Number of Tables
	Summary Read
	Summary Replicated Checks
	Summary Storage Buffer Rows
	Summary Uncompressed Cache Cells
	Summary Uptime
	Summary Write
Write Buffer	Summary Zookeeper Session
	Summary Zookeeper Watch
Write Buffer	Write Buffer File Descriptor Write

Table 8-41. Clickhouse Metrics (continued)

Component	Metrics
Replicated	Replicated Fetch
Memory	Memory Tracking
Query	Query Thread

Remote Check Metrics

Metrics are collected for object types such as HTTP, ICMP, TCP, and UDP.

HTTP Metrics

vRealize Operations Manager discovers metrics for HTTP remote checks.

HTTP Metrics

Table 8-42. HTTP Metrics

Metric Name	KPI
Availability	False
Response Code	False
Response Time	True
Result Code	False

ICMP Metrics

vRealize Operations Manager discovers metrics for the ICMP object type.

Table 8-43. ICMP Metrics

Metric Name	KPI
Availability	False
Average Response Time	True
Packet Loss (%)	False
Packets Received	False
Packets Transmitted	False
Result Code	False

TCP Metrics

vRealize Operations Manager discovers metrics for the TCP object type.

Table 8-44. TCP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

UDP Metrics

vRealize Operations Manager discovers metrics for the UDP object type.

Table 8-45. UDP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

Table 8-46. Virtual Machine Metrics

Metric Name	Description
Guest OS Services Total Number of Services	Number of out-of-the-box and user-defined services discovered in the VM.
Guest OS Services Number of User Defined Services	Number of user-defined services discovered in the VM.
Guest OS Services Number of OOTB Services	Number of out-of-the-box services discovered in the VM.
Guest OS Services Number of Outgoing Connections	Number of outgoing connection counts from the discovered services.
Guest OS Services Number of Incoming Connections	Number of incoming connection counts to the discovered services.

Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

Table 8-47. Service Summary Metrics

Metric Name	Description
Summary Incoming Connections Count	Number of incoming connections.
Summary Outgoing Connections Count	Number of outgoing connections.
Summary Connections Count	Number of incoming and outgoing connections.
Summary Pid	Process ID.

Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

Table 8-48. Service Performance Metrics

Metric Name	Description
Performance metrics group CPU	CPU usage in percentage.
Performance metrics group Memory	Memory usage in KB.
Performance metrics group IO Read Throughput	IO read throughput in KBps.
Performance metrics group IO Write Throughput	IO write throughput in KBps.

Service Type Metrics

Service discovery discovers metrics for service type objects.

Table 8-49. Service Type Metrics

Metric Name	Description
Number of instances	Number of instances of this service type.

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 Analytics Generated Metrics

The capacity engine computes and publishes metrics that can be found in the Capacity Analytics Generated group. These metrics help you to plan your resource use based on consumer demand.

Capacity Analytics Generated Metrics Group

Capacity analytics uses the capacity engine to analyze historical utilization and generate projected utilization. The engine takes the Demand and Usable Capacity (Total Capacity - HA - buffer) metrics as input and calculates the output metrics that belong to the capacity analytics generated metrics group.

The capacity analytics generated metrics group contains containers and each container contains three output metrics, which are Capacity Remaining, Recommended Size, and Recommended Total Capacity. It also contains the Capacity Remaining Percentage and Time Remaining metrics, which show the most constrained values of the containers.

For the capacity metrics group, full metric names include the name of the resource container. For example, if recommended size metrics are computed for CPU or memory, the actual metric names appear as `cpu|demand|recommendedSize` or `mem|demand|recommendedSize`.

Table 8-50. Capacity Metrics Group

Metric Name	Description
Time Remaining (Day(s))	The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: timeRemaining
Capacity Remaining	Capacity remaining is the maximum point between the usable capacity now and the projected utilization for 3 days into the future. If the projected utilization is above 100% of the usable capacity, Capacity Remaining is 0. Key: capacityRemaining
Capacity Remaining Percentage (%)	The percentage of Capacity Remaining of the most constrained resource with respect to the usable capacity. Key: capacityRemainingPercentage
Recommended Size	The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. Recommended Size excludes HA settings. Key: recommendedSize
Recommended Total Capacity	The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. Recommended Total Capacity excludes HA settings. Key: recommendedTotalCapacity

Capacity Analytics Generated Allocation Metrics

Capacity allocation metrics provide information about the allotment of capacity for Cluster Compute and Datastore Cluster Resources.

Metric Name	Description
Capacity Analytics Generated CPU Allocation Capacity Remaining (vCPUs)	For vSphere objects published on Cluster Compute Resource only. Capacity Remaining based on overcommit ratio (if configured in effective policy). Key: OnlineCapacityAnalytics cpu alloc capacityRemaining
Capacity Analytics Generated CPU Allocation Recommended Total Capacity (Cores)	For vSphere objects published on Cluster Compute Resource only. The recommended level of total capacity, to maintain a green state for time remaining for the given object. Key: OnlineCapacityAnalytics cpu alloc recommendedTotalSize
Capacity Analytics Generated CPU Allocation Time Remaining (Day(s))	For vSphere objects published on Cluster Compute Resource only. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out. Key: OnlineCapacityAnalytics cpu alloc timeRemaining
CPU Allocation Usable Capacity after HA and Buffer (vCPUs)	For vSphere objects published on Cluster Compute Resource only. The usable capacity (total capacity - HA) based on configured overcommit ratio. Key: cpu alloc usableCapacity
Capacity Analytics Generated CPU Allocation Recommended Size (Cores)	For vSphere objects published on Cluster Compute Resource only. The recommended level of usable capacity (total capacity - HA), to maintain a green state for time remaining for the given object. Key: OnlineCapacityAnalytics cpu alloc recommendedSize
vRealize Operations Manager Generated Properties CPU Allocation Overcommit Ratio Setting	For vSphere objects published on Cluster Compute Resource only. This property shows the allocation overcommit ratio for CPU provided in effective policy. Key: System Properties cpu alloc overcommitRatioSetting
vRealize Operations Manager Generated Properties CPU Allocation Buffer (%)	CPU buffer percent defined by policy setting for allocation based capacity computation. Key: Properties cpu alloc bufferSetting
Capacity Analytics Generated Memory Allocation Capacity Remaining (KB)	For vSphere objects published on Cluster Compute Resource only. Capacity Remaining based on overcommit ratio (if configured in effective policy). Key: OnlineCapacityAnalytics mem alloc capacityRemaining
Capacity Analytics Generated Memory Allocation Recommended Total Capacity (KB)	For vSphere objects published on Cluster Compute Resource only. The recommended level of total capacity, to maintain a green state for time remaining for the given object. Key: OnlineCapacityAnalytics mem alloc recommendedTotalSize

Metric Name	Description
Capacity Analytics Generated Memory Allocation Time Remaining (Day(s))	For vSphere objects published on Cluster Compute Resource only. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out. Key: OnlineCapacityAnalytics mem alloc timeRemaining
Memory Allocation Usable Capacity (KB)	For vSphere objects published on Cluster Compute Resource only. The usable capacity (total capacity - HA) based on configured overcommit ratio. Key: mem alloc usableCapacity
Capacity Analytics Generated Memory Allocation Recommended Size (KB)	For vSphere objects published on Cluster Compute Resource only. The recommended level of usable capacity (total capacity - HA), to maintain a green state for time remaining for the given object. Key: OnlineCapacityAnalytics mem alloc recommendedSize
vRealize Operations Manager Generated Properties Memory Allocation Overcommit Ratio Setting	For vSphere objects published on Cluster Compute Resource only. This property shows the allocation overcommit ratio for Memory provided in effective policy. Key: System Properties mem alloc overcommitRatioSetting
vRealize Operations Manager Generated Properties Memory Allocation Buffer (%)	Memory buffer percent defined by policy setting for allocation based capacity computation. Key: System Properties mem alloc bufferSetting
Capacity Analytics Generated Disk Space Allocation Capacity Remaining (GB)	For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. Capacity Remaining based on overcommit ratio (if configured in effective policy). Key: OnlineCapacityAnalytics diskspace alloc capacityRemaining
Capacity Analytics Generated Disk Space Allocation Recommended Size (GB)	For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. The recommended level of total capacity to maintain a green state for time remaining for the given object. Key: OnlineCapacityAnalytics diskspace alloc recommendedSize
Capacity Analytics Generated Disk Space Allocation Time Remaining (Day(s))	For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out. Key: OnlineCapacityAnalytics diskspace alloc timeRemaining
Disk Space Allocation Usable Capacity (GB)	For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. Usable capacity based on overcommit ratio (if configured in effective policy). Key: diskspace alloc usableCapacity

Metric Name	Description
vRealize Operations Manager Generated Properties Disk Space Allocation Overcommit Ratio Setting	For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. This property shows the allocation overcommit ratio for Disk Space provided in effective policy. key: System Properties diskspace alloc overcommitRatioSetting
vRealize Operations Manager Generated Properties Disk Space Allocation Buffer (%)	Disk Space buffer percent defined by policy setting for allocation based capacity computation. Key: System Properties diskspace alloc bufferSetting

Capacity Analytics Generated Profiles Metrics

Profiles metrics provide information about the profile specific capacity for Cluster Compute, Datastore Cluster, Data Center, Custom Data Center, and vCenter Server resources.

Metric Name	Description
Capacity Analytics Generated Capacity Remaining (Profile)	Published on Cluster Compute Resource. Calculated as a minimum of all Profiles capacityRemainingProfile_<profile uuid> metrics. Key: OnlineCapacityAnalytics capacityRemainingProfile
Capacity Analytics Generated Capacity Remaining (Profile)	Published on Datastore Cluster Resource. Calculated as a minimum of all Profiles capacityRemainingProfile_<profile uuid> metrics. Key: OnlineCapacityAnalytics capacityRemainingProfile
Capacity Analytics Generated Capacity Remaining (Profile)	Published on Data Center, Custom Data Center and vCenter Server Resources. Computed as a sum of OnlineCapacityAnalytics capacityRemainingProfile metric of descendant Cluster Compute Resources. Key: OnlineCapacityAnalytics capacityRemainingProfile

Capacity Demand Model Metrics

Demand model metrics provide information about the usable capacity and projected utilization of resources across VMs, Host Systems, Cluster Compute, Datastore Cluster, Data Center, Custom Data Center, and vCenter Server resources.

Metric Name	Description
Capacity Analytics Generated CPU Capacity Remaining (MHz)	Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days. Key: OnlineCapacityAnalytics cpu capacityRemaining
Capacity Analytics Generated CPU Recommended Size (MHz)	Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics cpu recommendedSize

Metric Name	Description
Capacity Analytics Generated CPU Time Remaining (Day(s))	Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics cpu timeRemaining
Capacity Analytics Generated Disk Space Capacity Remaining (GB)	Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace capacityRemaining
Capacity Analytics Generated Disk Space Recommended Size (GB)	Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics diskspace recommendedSize
Capacity Analytics Generated Disk Space Time Remaining (Day(s))	Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace timeRemaining
Capacity Analytics Generated Memory Capacity Remaining (KB)	Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics mem capacityRemaining
Capacity Analytics Generated Memory Recommended Size (KB)	Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics mem recommendedSize
Capacity Analytics Generated Memory Time Remaining (Day(s))	Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics mem timeRemaining
Capacity Analytics Generated CPU Demand Capacity Remaining (MHz)	Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics cpu demand capacityRemaining
vRealize Operations Manager Generated Properties CPU Demand Buffer (%)	CPU buffer percent defined by policy setting for demand based capacity computation. Key: System Properties cpuldemand bufferSetting
Capacity Analytics Generated CPU Demand Recommended Size (MHz)	Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics cpuldemand recommendedSize

Metric Name	Description
Capacity Analytics Generated CPU Demand Time Remaining (Day(s))	Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics cpu demand timeRemaining
Capacity Analytics Generated Disk Space Demand Capacity Remaining (GB)	Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace demand capacityRemaining
vRealize Operations Manager Generated Properties Disk Space Demand Buffer (%)	Disk Space buffer percent defined by policy setting for demand based capacity computation. System Properties diskspace demand bufferSetting
Capacity Analytics Generated Disk Space Demand Recommended Size (GB)	Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics diskspace demand recommendedSize
Capacity Analytics Generated Disk Space Demand Time Remaining (Day(s))	Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace demand timeRemaining
Capacity Analytics Generated Memory Demand Capacity Remaining (KB)	Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics mem demand capacityRemaining
vRealize Operations Manager Generated Properties Memory Demand Buffer (%)	Memory buffer percent defined by policy setting for demand based capacity computation. Key: System Properties mem demand bufferSetting
Capacity Analytics Generated Memory Demand Recommended Size (KB)	Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics mem demand recommendedSize
Capacity Analytics Generated Memory Demand Time Remaining (Day(s))	Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics mem demand timeRemaining
Capacity Analytics Generated Disk Space Usage Capacity Remaining (GB)	Published on Datastore. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace total capacityRemaining

Metric Name	Description
Capacity Analytics Generated Disk Space Usage Recommended Size (GB)	Published on Datastore. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics diskspace total recommendedSize
Capacity Analytics Generated Disk Space Usage Time Remaining (Day(s))	Published on Datastore. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace total timeRemaining
Capacity Analytics Generated CPU Demand Capacity Remaining (MHz)	Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics cpu demand capacityRemaining
Capacity Analytics Generated CPU Demand Recommended Size (MHz)	Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time. Key: OnlineCapacityAnalytics cpu demand recommendedSize
Capacity Analytics Generated CPU Demand Recommended Total Capacity (MHz)	Published on Cluster Compute Resource. The recommended level of total capacity to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics cpu demand recommendedTotalSize
Capacity Analytics Generated CPU Demand Time Remaining (Day(s))	Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics cpu demand timeRemaining
Capacity Analytics Generated Disk Space Demand Capacity Remaining (GB)	Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace demand capacityRemaining
Capacity Analytics Generated Disk Space Demand Recommended Size (GB)	Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics diskspace demand recommendedSize
Capacity Analytics Generated Disk Space Demand Time Remaining (Day(s))	Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace demand timeRemaining

Metric Name	Description
Capacity Analytics Generated Memory Demand Capacity Remaining (KB)	Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics mem demand capacityRemaining
Capacity Analytics Generated Memory Demand Recommended Size (KB)	Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics mem demand recommendedSize
Capacity Analytics Generated Memory Demand Recommended Total Capacity (KB)	Published on Cluster Compute Resource. The recommended level of total capacity to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics mem demand recommendedTotalSize
Capacity Analytics Generated Memory Demand Time Remaining (Day(s))	Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics mem demand timeRemaining
Capacity Analytics Generated Disk Space Usage Capacity Remaining (GB)	Published on Datastore Cluster. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace total capacityRemaining
Capacity Analytics Generated Disk Space Usage Recommended Size (GB)	Published on Datastore Cluster. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics diskspace total recommendedSize
Capacity Analytics Generated Disk Space Usage Time Remaining (Day(s))	Published on Datastore Cluster. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace total timeRemaining
Capacity Analytics Generated CPU Demand Capacity Remaining (MHz)	Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics cpu demand capacityRemaining
Capacity Analytics Generated CPU Demand Recommended Size (MHz)	Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics cpu demand recommendedSize

Metric Name	Description
Capacity Analytics Generated CPU Demand Recommended Total Capacity (MHz)	Published on Datacenter, Custom Datacenter, vCenter. The recommended level of total capacity to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics cpu demand recommendedTotalSize
Capacity Analytics Generated CPU Demand Time Remaining (Day(s))	Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics cpu demand timeRemaining
Capacity Analytics Generated Disk Space Demand Capacity Remaining (GB)	Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics diskspace demand capacityRemaining
Capacity Analytics Generated Disk Space Demand Recommended Size (GB)	Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics diskspace demand recommendedSize
Capacity Analytics Generated Disk Space Demand Time Remaining (Day(s))	Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics diskspace demand timeRemaining
Capacity Analytics Generated Memory Demand Capacity Remaining (KB)	Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future. Key: OnlineCapacityAnalytics mem demand capacityRemaining
Capacity Analytics Generated Memory Demand Recommended Size (KB)	Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics mem demand recommendedSize
Capacity Analytics Generated Memory Demand Recommended Total Capacity (KB)	Published on Datacenter, Custom Datacenter, vCenter. The recommended level of total capacity to maintain a green state for the time remaining. Key: OnlineCapacityAnalytics mem demand recommendedTotalSize
Capacity Analytics Generated Memory Demand Time Remaining (Day(s))	Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity. Key: OnlineCapacityAnalytics mem demand timeRemaining

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 8-51. Badge Metrics

Metric Name	Description
Badge Compliance	Overall score for compliance, on a scale of 100.
Badge Efficiency	Overall score for efficiency. The final score is between 1-100. Where Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1. The score is derived from the criticality of alerts in the Efficiency category.
Badge Health	Overall score for health. The final score is between 1-100. Where Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1. The score is derived from the criticality of alerts in the Health category.
Badge Risk	Overall score for risk. The final score is between 1-100. Where Green - 0, Yellow - 25, Orange - 50, Red - 75, Unknown: -1. The score is derived from the criticality of alerts in the Risk category.

System Metrics

System metrics provide information used to monitor the health of the system. They help you to identify problems in your environment.

Table 8-52. System Metrics

Metric Name	Description
vRealize Operations Generated Self - Health Score	This metric displays the system health score of self resource. The value ranges from 0 to 100 depending on noise and the number of alarms. Key: System Attributes health
vRealize Operations Generated Self - Metric Count	This metric displays the number of metrics that the adapter generates for the given object. This value does not include the number of metrics generated by vRealize Operations Manager , such as, Badge metrics, vRealize Operations Generated metrics and metrics generated by Capacity Engine Key: System Attributes all_metrics
vRealize Operations Generated Total Anomalies	This metric displays the number of active anomalies (symptoms, events, DT violations) on the object and its children. In previous versions of vRealize Operations Manager, this metric used to be named vRealize Operations Generated Self - Total Anomalies. Key: System Attributes total_alarms

Table 8-52. System Metrics (continued)

Metric Name	Description
vRealize Operations Generated Full Set - Metric Count	This metric displays the number of metrics that the adapter of the children of the given object generates. Key: System Attributes child_all_metrics
vRealize Operations Generated Availability	This metric value is computed based on the adapter instance statuses monitoring the resource. Resource availability is displayed as 0-down, 1-Up, -1-Unknown. Key: System Attributes availability
vRealize Operations Generated Alert Count Critical	This metric displays the number of critical alerts on the object and its children. Key: System Attributes alert_count_critical
vRealize Operations Generated Alert Count Immediate	This metric displays the number of immediate alerts on the object and its children. Key: System Attributes alert_count_immediate
vRealize Operations Generated Alert Count Warning	This metric displays the number of active warning alerts on the object and its children. Key: System Attributes alert_count_warning
vRealize Operations Generated Alert Count Info	This metric displays the number of active info alerts on the object and its children. Key: System Attributes alert_count_info
vRealize Operations Generated Total Alert Count	This metric displays the sum of all alert count metrics. In previous versions of vRealize Operations Manager, this metric was named vRealize Operations Generated Full Set - Alert Count. Key: System Attributes total_alert_count
vRealize Operations Generated Self-Alert Count	This metric displays the number of all alerts on the object. Key: System Attributes self_alert_count

Log Insight Generated Metrics

The metrics in the Log Insight Generated group provide information that you can use to observe or troubleshoot vRealize Operations Manager for failures and to monitor performance.

When vRealize Operations Manager is integrated with Log Insight and metric calculation is enabled, Log Insight calculates the number of logs corresponding to different queries and sends them as metrics to vRealize Operations Manager. These metrics are calculated for vCenter objects, host objects, and virtual machine objects. The metrics can be mapped to a vRealize Operations Manager object based on the Log Insight field *vmw_vrops_id*, which is constructed based on hostname or source fields.

Table 8-53. Log Insight Generated Metrics

Metric Name	Description
Log Insight Generated Error Count	The number of error logs for the selected object. Key: log_insight_generated error_count
Log Insight Generated Total Log Count	The total number of logs for the selected object. Key: log_insight_generated total_log_count
Log Insight Generated Warning Count	The number of warning logs for the selected object. Key: log_insight_generated warning_count

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 8-54. 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.
LocalFDAItSize	Number of alternative forward data entries	Number of locally stored primary and redundant entries in alternative forward data region.

Table 8-54. Analytics Metrics (continued)

Metric Key	Metric Name	Description
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
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 8-55. 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)

Table 8-55. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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
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

Table 8-55. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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
UpdateResourceCache Duration MaxDuration	Maximum	Maximum
UpdateResourceCache ModificationEstimateCount TotalCount	Total	The number of estimated modifications done during each resource cache object update.
UpdateResourceCache ModificationEstimateCount AvgCount	Average	Average
UpdateResourceCache ModificationEstimateCount MinCount	Minimal	Minimal
UpdateResourceCache ModificationEstimateCount 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

Table 8-55. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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 8-56. 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)
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

Table 8-56. Dynamic Threshold Calculation Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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 8-57. 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 8-58. 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
RejectedFDAItCount	Number of rejected alternative forward data	Number of rejected alternative forward data
SentFDCount	Number of sent objects	Number of sent objects
SentFDAItCount	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.

Table 8-58. Collector Metrics (continued)

Metric Key	Metric Name	Description
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 8-59. 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 8-60. 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 8-61. 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)

Table 8-61. Storage Thread Pool Metrics for FSDB (continued)

Metric Key	Metric Name	Description
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 8-62. 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 8-63. 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)

Table 8-63. API Call Metrics for the Product UI (continued)

Metric Key	Metric Name	Description
APICalls AlertRequestCount	Alert request count	Alert request count
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 8-64. 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 8-65. 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 vRealize Operations Manager API objects.

Table 8-66. 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 8-67. 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 8-68. 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
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 8-69. 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 8-70. 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)

Table 8-70. API Call Metrics for Cluster and Slice Administration (continued)

Metric Key	Metric Name	Description
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 8-71. Watchdog Metrics

Metric Key	Metric Name	Description
ServiceCount	Service Count	Service Count

Service Metrics

Service metrics provide information about watchdog activity.

Table 8-72. 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 8-73. 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 8-74. 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 8-75. 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 8-76. 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

Table 8-76. Resource Limit Metrics for the Node (continued)

Metric Key	Metric Name	Description
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 8-77. 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 state 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
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 8-78. 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

Table 8-78. Network Interface Metrics for the Node (continued)

Metric Key	Metric Name	Description
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 8-79. 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 8-80. 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 8-81. 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 8-82. 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 8-83. 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)

Table 8-83. CPU Metrics for the Node (continued)

Metric Key	Metric Name	Description
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 8-84. 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 8-85. 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 8-86. 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 did not respond.

Table 8-87. 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 8-88. 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 8-89. 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

Table 8-89. DT Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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 8-90. 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
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 8-91. 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

Table 8-91. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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
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)

Table 8-91. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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
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 CommittedMemory	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

Table 8-91. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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
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 8-92. 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 8-93. 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 8-94. 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 8-95. 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 8-96. 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 8-97. Activity Metrics for Persistence

Metric Key	Metric Name	Description
Activity RunningCount	Number Running	Number Running
Activity ExecutedCount	Number Executed	Number Executed
Activity SucceededCount	Number Succeeded	Number Succeeded
Activity FailedCount	Number Failed	Number Failed

Controller XDB Metrics

Controller metrics relate to the primary database.

Table 8-98. 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 8-99. 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)

Table 8-99. Alarm XDB Metrics for Persistence (continued)

Metric Key	Metric Name	Description
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 8-100. 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 8-100. 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 stopped

Remote Collector Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager remote collector node objects.

Table 8-101. Remote Collector Metrics

Metric Key	Metric Name	Description
ComponentCount	Component Count	The number of vRealize Operations Manager Objects reporting for this node.

Table 8-102. Memory Metrics for the Remote Collector

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 8-103. 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 8-104. 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 8-105. 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

Table 8-105. Network Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
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
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 8-106. 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

Table 8-106. Network Interface Metrics for the Remote Collector (continued)

Metric Key	Metric Name	Description
net iface txErrors	Transmit packet errors	Number of transmit errors
net iface txOverruns	Transmit packet overruns	Number of transmit overruns

Table 8-107. 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
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 8-108. 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 8-109. 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 8-110. 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 8-111. 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)
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 8-112. 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 8-113. 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 8-114. 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.

vRealize Automation 8.x Metrics

vRealize Automation 8.x collects metrics for objects such as, cloud zone, project, deployment, blueprint, cloud account, user, and cloud automation services world Instance.

Blueprint Metrics

vRealize Automation 8.x collects metrics for objects such as blueprint object.

Table 8-115. Blueprint Metrics

Property Name	Metrics
Summary	VMCount

Project Metrics

vRealize Automation 8.x collects metrics for objects such as project object.

Table 8-116. Project Metrics

Property Name	Metrics
Summary	VMCount
Summary	TotalDeployments
Summary	TotalCloudZones
Summary	TotalBlueprints
Summary	Metering Additional price
Summary	Metering CPU Price
Summary	Metering Memory price
Summary	Metering Storage Price
Summary	Metering Total price

Deployment Metrics

vRealize Automation 8.x collects the metrics for the deployment object.

Table 8-117. Deployment Metrics

Property Name	Metrics
Summary	Metering Additional price
Summary	Metering CPU Price
Summary	Metering Memory price
Summary	Metering Storage Price
Summary	Metering Total price
Summary	Metering Partial price

Organization Metrics

vRealize Automation 8.x collects the metrics for the organization object.

Table 8-118. Organization Metrics

Property Name	Metrics
Summary	TotalBlueprints
Summary	TotalProjects
Summary	VMCount

Table 8-118. Organization Metrics (continued)

Property Name	Metrics
Summary	TotalDeployments
Summary	TotalCloudZones

vRealize Adapter 8.x Metrics

vRealize Automation 8.x collects the metrics for the vRealize adapter object.

Table 8-119. vRealize Adapter 8.x Metrics

Property Name	Metrics
Summary	TotalCloudZones
Summary	VMCount
Summary	TotalDeployments
Summary	TotalBlueprints
Summary	TotalProjects

Cloud Automation Services World Metrics

vRealize Automation 8.x collects the metrics for the Cloud Automation Services world object.

Table 8-120. Cloud Automation Services World Metrics

Property Name	Metrics
Summary	TotalDeployments
Summary	VMCount
Summary	TotalCloudZones
Summary	TotalProjects
Summary	TotalBlueprints

Cloud Automation Services Entity Status Metrics

vRealize Automation 8.x collects the metrics for the Cloud Automation Services (CAS) entity status object.

Table 8-121. Cloud Automation Services Entity Status Metrics

Property Name	Metrics
Summary	TotalClusters

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 **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|Total Bus Resets
- Disk I/O|Total Commands Aborted per second

The following Disk I/O metrics are disabled by default:

- 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 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|Reads Per Second (IOPS)
- Read Cache|Read Latency (ms)
- Read Cache|Writes Per Second (IOPS)
- Read Cache|Write Latency (ms)

The following read cache metrics are disabled by default:

- Read Cache|Read I/O Count
- 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|Reads Per Second (IOPS)
- Write Buffer|Read Latency (ms)
- Write Buffer|Writes Per Second (IOPS)
- Write Buffer|Write Latency (ms)

The following write buffer metrics are disabled by default:

- Write Buffer|Read I/O Count
- 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

Cache De-stage Metrics for vSAN Disk Groups

The vRealize Operations Manager collects cache de-stage metrics for the vSAN disk groups.

Cache de-stage metrics include:

- Bytes De-stage from SSD
- Zero-bytes De-stage

Resync Traffic Metrics for vSAN Disk Groups

The vRealize Operations Manager collects resync traffic metrics for the vSAN disk groups.

Resync traffic metrics include:

- Read IOPS for Resync Traffic
- Write IOPS for Resync Traffic
- Read Throughput for Resync Traffic
- Write Throughput for Resync Traffic
- Read Latency for Resync Traffic
- Write Latency for Resync Traffic

Metrics for vSAN Cluster

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN cluster.

vRealize Operations Manager enhances the capacity calculation for vSAN, using the new slack space provided by the new vSAN API. Cost calculation is still done using the old way which reserves 30% memory for Slack Overhead.

Metrics for vSAN cluster include:

Component	Metrics
Component Limit	<ul style="list-style-type: none"> ■ vSAN Component Limit Component Limit Used (%) ■ vSAN Component Limit Total Component Limit ■ vSAN Component Limit Used Component Limit
Disk Space	<ul style="list-style-type: none"> ■ vSAN Disk Space Disk Space Used (%) ■ vSAN Disk Space Total Disk Space (GB) ■ vSAN Disk Space Used Disk Space (GB) ■ vSAN Disk Space Usable Capacity (GB)
Read Cache	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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 ■ Summary Total vSAN CPU Used (MHz) ■ Summary Max vSAN CPU Ready ■ Summary Worst VM Disk Latency
KPI	<ul style="list-style-type: none"> ■ KPI Sum Host VMKernel Packets Dropped ■ KPI Count Disk Group Congestion Above 50 ■ KPI Max Disk Group Congestion ■ KPI Sum Disk Group Errors ■ KPI Min Disk Group Capacity Free ■ KPI Min Disk Group Read Cache Hit Rate ■ KPI Min Disk Group Write Buffer Free ■ KPI Max Disk Group Read Cache/Write Buffer Latency ■ KPI Max Capacity Disk Latency ■ KPI Max Capacity Disk IOPS
IO Size	<ul style="list-style-type: none"> ■ vSAN Performance I/O Size (KB) ■ vSAN Performance Read I/O Size (KB) ■ vSAN Performance Write I/O Size (KB)

Component	Metrics
Resynchronization Status (Metrics applicable for vSAN 6.7 and later)	<ul style="list-style-type: none"> ■ vSAN Resync Bytes left to resync (bytes) ■ vSAN Resync Resyncing Objects
Stretched Cluster	<ul style="list-style-type: none"> ■ vSAN Stretched Cluster Latency Between Sites Preferred and Secondary (ms) ■ vSAN Stretched Cluster Latency Between Sites Preferred and Witness (ms) ■ vSAN Stretched Cluster Latency Between Sites Secondary and Witness (ms)
File Share	<ul style="list-style-type: none"> ■ vSAN FileServices totalShareCount
File Service	<ul style="list-style-type: none"> ■ vSAN File Services File Shares Used Disk Space (GB) ■ vSAN File Services Root FS Used Disk Space (GB) ■ vSAN File Services File Shares Count
Slack Space	<ul style="list-style-type: none"> ■ vSAN Slack Space Internal Operations Capacity (GB) ■ vSAN Slack Space Host Rebuild Capacity (GB) ■ vSAN Slack Space Transient Capacity Used (GB)

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 a vSAN enabled host include:

Component	Metrics
Component Limit	<ul style="list-style-type: none"> ■ vSAN Component Limit Component Limit Used (%) ■ vSAN Component Limit Total Component Limit ■ vSAN Component Limit Used Component Limit
Disk Space	<ul style="list-style-type: none"> ■ vSAN Disk Space Disk Space Used (%) ■ vSAN Disk Space Total Disk Space (GB) ■ vSAN Disk Space Used Disk Space (GB)
Read Cache	<ul style="list-style-type: none"> ■ vSAN Read Cache Read Cache Reserved (%) ■ vSAN Read Cache Reserved Read Cache Size (GB) ■ vSAN Read Cache Total Read Cache Size (GB)
Performance Metrics	
<ul style="list-style-type: none"> ■ Network 	<ul style="list-style-type: none"> ■ vSAN Performance Network Inbound Packets Loss Rate ■ vSAN Performance Network Outbound Packets Loss Rate ■ vSAN Performance Network <vnic> Inbound Packets Loss rate (%) ■ vSAN Performance Network <vnic> Outbound Packets Loss Rate (%) ■ vSAN Performance Network <vnic> Inbound Packets Per second ■ vSAN Performance Network <vnic> Outbound Packets Per second ■ vSAN Performance Network <vnic> Throughput Inbound (KBps) ■ vSAN Performance Network <vnic> Throughput Outbound (KBps)

Component	Metrics
■ CPU Utilization	<ul style="list-style-type: none"> ■ vSAN Performance CPU Ready (%) ■ vSAN Performance CPU Usage (%) ■ vSAN Performance CPU Used (MHz) ■ vSAN Performance CPU Core Utilization (%) (For Hyper-Threading Technology)
■ PCPU Utilization	<ul style="list-style-type: none"> ■ vSAN Performance PCPU Ready (%) ■ vSAN Performance CPU PCPU Usage (%)
■ Memory	<ul style="list-style-type: none"> ■ vSAN Performance Memory Usage (%) ■ vSAN Performance Memory Used (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
- Capacity | Usable Capacity

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	<ul style="list-style-type: none"> ■ Performance Bus Resets ■ Performance Commands Aborted Per Second <p>The following performance metrics are disabled by default:</p> <ul style="list-style-type: none"> ■ 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 <hr/> Note SMART data collection is disabled by default. To enable SMART data collection, ensure that the <code>Enable SMART data collection</code> instance identifier is set to <code>true</code> . For proper data collection, ensure that ESXi hosts in your vCenter Server inventory have CIM service enabled and CIM providers for each SMART metric installed.	<ul style="list-style-type: none"> ■ 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 ■ SCSI SMART Statistics Worst Media Wearout Indicator ■ SCSI SMART Statistics Worst Write Error Count ■ SCSI SMART Statistics Worst Read Error Count ■ SCSI SMART Statistics Worst Power-on Hours ■ SCSI SMART Statistics Power Cycle Count ■ SCSI SMART Statistics Worst Power Cycle Count ■ SCSI SMART Statistics Worst Reallocated Sector Count ■ SCSI SMART Statistics Worst Raw Read Error Rate ■ SCSI SMART Statistics Worst Driver Rated Max Temperature ■ SCSI SMART Statistics Worst Write Sectors TOT Count ■ SCSI SMART Statistics Worst Read Sectors TOT Count ■ SCSI SMART Statistics Worst Initial Bad Block Count
Capacity	<ul style="list-style-type: none"> ■ vSAN Health Capacity Total Disk Capacity (GB) ■ vSAN Health Capacity Used Disk Capacity (GB)

Component	Metrics
Congestion Health	■ vSAN Health Congestion Health Congestion Value
Performance	<ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ Performance Bus Resets ■ Performance Commands Aborted Per Second <p>The following performance metrics are disabled by default:</p> <ul style="list-style-type: none"> ■ ■ 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 <hr/> Note SMART data collection is disabled by default. To enable SMART data collection, ensure that the <code>Enable SMART data collection</code> instance identifier is set to true. For proper data collection, ensure that ESXi hosts in your vCenter Server inventory have CIM service enabled and CIM providers for each SMART metric installed.	<ul style="list-style-type: none"> ■ 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 ■ SCSI SMART Statistics Worst Media Wearout Indicator ■ SCSI SMART Statistics Worst Write Error Count ■ SCSI SMART Statistics Worst Read Error Count ■ SCSI SMART Statistics Worst Power-on Hours ■ SCSI SMART Statistics Power Cycle Count ■ SCSI SMART Statistics Worst Power Cycle Count ■ SCSI SMART Statistics Worst Reallocated Sector Count ■ SCSI SMART Statistics Worst Raw Read Error Rate ■ SCSI SMART Statistics Worst Driver Rated Max Temperature ■ SCSI SMART Statistics Worst Write Sectors TOT Count ■ SCSI SMART Statistics Worst Read Sectors TOT Count ■ SCSI SMART Statistics Worst Initial Bad Block Count
Capacity	<ul style="list-style-type: none"> ■ vSAN Health Total Disk Capacity (GB) ■ vSAN Health Used Disk Capacity (GB) ■ vSAN FileServices FileSharesUsedDiskSpace ■ vSAN FileServices RootFsUsedDiskSpace

Component	Metrics
Congestion Health	vSAN Health Congestion Value
Performance	<ul style="list-style-type: none"> ■ 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 ■ vSAN Performance vSAN Layer Total IOPS

Properties for vSAN capacity disk include:

- Name
- Size
- Vendor
- Type
- Queue Depth

Metrics for vSAN Fault Domain Resource Kind

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN stretched cluster with fault domain.

Metrics for vSAN fault domain resource kind includes:

- CPU
 - Demand
 - Demand (MHz)
 - Demand without overhead (MHz)
 - Overhead (MHz)
 - Reserved Capacity (MHz)
 - Total Capacity (MHz)

- VM CPU Usage (MHz)
 - Workload (%)
- Disk Space
 - Demand
 - Workload (%)
- Memory
 - Contention (KB)
 - Demand
 - Host Usage (KB)
 - Machine Demand (KB)
 - Reserved Capacity (KB)
 - Total Capacity (KB)
 - Utilization (KB)
 - Workload (%)
- vSAN
 - Disk Space
 - Total Disk Space (GB)
 - Used Disk Space (GB)

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 vSAN File Server

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN File Server.

Metrics for vSAN File Server

Component	Metrics
File Server	<ul style="list-style-type: none"> ■ vSAN Disk Space File Shares Used Disk Space (GB) ■ vSAN Summary File Shares Count

Metrics for vSAN File Share

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN File Share.

Metrics for vSAN File Share

Component	Metrics
Disk Space	<ul style="list-style-type: none"> ■ vSAN Disk Space Used Disk Space (GB)
Read Performance	<ul style="list-style-type: none"> ■ vSAN Performance Read Throughput Requested (MBps) ■ vSAN Performance Read Throughput Transferred (MBps) ■ vSAN Performance Read IOPS ■ vSAN Performance Read Latency (ms)
Write Performance	<ul style="list-style-type: none"> ■ vSAN Performance Write Throughput Requested (MBps) ■ vSAN Performance Write Throughput Transferred (MBps) ■ vSAN Performance Write IOPS ■ vSAN Performance Write Latency (ms)

Capacity Model for vSAN Objects

The capacity model introduced in vRealize Operations Manager 6.7 now extends the support for vSAN objects like, vSAN cluster, Fault domains, and Cache/Capacity disks. The Capacity tab provides Time Remaining data for the selected vSAN cluster, Fault domain, Cache/Capacity Disk objects. The information is presented in a graphical format.

Where You Find the Capacity Tab

In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. The Object details page appears. Click the **Capacity** tab.

The vRealize Operations Manager defines the capacity model for the following vSAN resource containers:

- vSAN Cluster
 - Disk Space
- vSAN Fault Domain
 - CPU
 - Memory
 - Disk Space
- vSAN Cache/Capacity Disk
 - Disk Space

Understanding the Capacity Tab

For the selected vSAN resource, the capacity tab lists the capacity used and Time Remaining until the associated CPU, memory, and disk space resources, respectively, run out.

- If you select the vSAN cluster, the capacity tab lists the capacity used and time remaining until the associated disk space runs out.
- If you select the vSAN Fault Domain, the capacity tab lists the capacity used and time remaining until the associated CPU, memory, and disk space resources run out.
- If you select the vSAN Cache/Capacity Disk Space, the capacity tab lists capacity used and time remaining until the associated disk space runs out.

The available graph depicts - for your choice of CPU, memory, or disk space - the amount of resource used, plotted against time. A line on the graph shows 100 percent usable capacity and a trend line projects how swiftly resource use is approaching 100 percent. The time line shows when the selected resource is to reach capacity.

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 as Linux, AIX, Solaris, and Windows. The Operating Systems plug-in also collects metrics for Windows services, Script services, and Multiprocess services.

End Point Operations Management agents discover file systems and automatically monitor them for read/write rates, total capacity, used capacity, and so on.

AIX Metrics

The Operating Systems Plug-in discovers the metrics for the AIX object type. AIX 6.1 and 7.1 are supported.

Table 8-122. 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
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

Table 8-122. AIX Metrics (continued)

Name	Category	KPI
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
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

Table 8-122. AIX Metrics (continued)

Name	Category	KPI
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

Table 8-122. AIX Metrics (continued)

Name	Category	KPI
Number of CPUs	UTILIZATION	False
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

Table 8-122. AIX Metrics (continued)

Name	Category	KPI
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 8-123. 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

Table 8-123. Linux Metrics (continued)

Name	Category	KPI
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
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 SoftIrq 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 SoftIrq Time per Minute	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Irq Time	UTILIZATION	False

Table 8-123. Linux Metrics (continued)

Name	Category	KPI
Cpu SoftIrq	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
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

Table 8-123. Linux Metrics (continued)

Name	Category	KPI
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
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False

Table 8-123. Linux Metrics (continued)

Name	Category	KPI
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

Table 8-123. Linux Metrics (continued)

Name	Category	KPI
Used Memory (- buffers/cache)	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
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 8-124. 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

Table 8-124. Solaris Metrics (continued)

Name	Category	KPI
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
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

Table 8-124. Solaris Metrics (continued)

Name	Category	KPI
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
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

Table 8-124. Solaris Metrics (continued)

Name	Category	KPI
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
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

Table 8-124. Solaris Metrics (continued)

Name	Category	KPI
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 8-125. Microsoft Windows Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False

Table 8-125. Microsoft Windows Metrics (continued)

Name	Category	KPI
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

Table 8-125. Microsoft Windows Metrics (continued)

Name	Category	KPI
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 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

Table 8-125. Microsoft Windows Metrics (continued)

Name	Category	KPI
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
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

Table 8-125. Microsoft Windows Metrics (continued)

Name	Category	KPI
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 8-126. 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

Table 8-126. Windows Services Metrics (continued)

Name	Category	KPI
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. The metrics will be available only if the shell script is configured.

Table 8-127. Script Metrics

Name	Category	KPI	Description
Resource Availability	AVAILABILITY	True	Displays if the script is available or not. If the value is "0" the script is unavailable. If the value is "100" the script it available. Key: Availability Resource Availability
Execution Time	THROUGHPUT	True	Time spent to run the script. Key: Throughput Execution Time (ms)
Result Value	UTILIZATION	True	Exit value of the script. If the script contains "echo 1", the the value is 1. If the script contains "echo 0", the value will be 0. Key: Utilization Result value

Multiprocess Service Metrics

The Operating Systems Plug-in discovers the metrics for the Multiprocess service.

Table 8-128. 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 8-129. HTTP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Last Modified	AVAILABILITY	False

Table 8-129. HTTP Check Metrics (continued)

Name	Category	KPI
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
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 8-130. 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 8-131. 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
State LAST_ACK	THROUGHPUT	False
State CLOSING	THROUGHPUT	False
All Outbound Connections	THROUGHPUT	False
State SYN_RECV	THROUGHPUT	False
State FIN_WAIT1	THROUGHPUT	False

Metrics for Microsoft Azure

vRealize Operations Manager collects metrics for Microsoft Azure adapter objects.

On the menu, click **Environment > All Objects > Microsoft Azure Adapter** and expand an object. Select one of the object instances and click the **Metrics** tab.

Virtual Machine Metrics

The following metrics are available for each Virtual Machine instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
Percentage CPU	Metric	Percent	Average	The percentage of allocated compute units that are currently in use by the Virtual Machine.
OS Type	Property	String	Not applicable.	The type of operating system.
OS VHD URI	Property	String	Not applicable.	The virtual hard disk URI of the operating system.
Service Tier	Property	String	Not applicable.	The size of the Virtual Machine.
FQDN	Property	String	Not applicable.	The fully qualified domain name of the Virtual Machine.
Disk Read Bytes	Metric	Bytes	Average	The average bytes read from the disk during the monitoring period.
Disk Write Bytes	Metric	Bytes	Average	The average bytes written to the disk during the monitoring period.
Disk Read Operations/Sec	Metric	Count Per Second	Average	The average number of requests read from the disk per second.
Disk Write Operations/Sec	Metric	Count Per Second	Average	The average number of requests written to the disk per second.
Network In Total	Metric	Bytes	Total	The number of bytes received on all network interfaces by the Virtual Machine.
Network Out Total	Metric	Bytes	Total	The number of bytes out on all network interfaces by the Virtual Machine.

Cosmos DB Metrics

The following metrics are available for each Cosmos DB instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/cosmos-db/cosmos-db-azure-monitor-metrics>.

Name	Type	Unit	Aggregation Type	Description
Available Storage	Metric	Bytes	Total	The total available storage reported at 5-minutes granularity per region.
Data Usage	Metric	Bytes	Total	The total data usage reported at 5-minutes granularity per region.
Document Count	Metric	Count	Total	The total document count reported at 5-minutes granularity per region.
Document Quota	Metric	Bytes	Total	The total storage quota reported at 5-minutes granularity per region.
Index Usage	Metric	Bytes	Total	The total index usage reported at 5-minutes granularity per region.

SQL Server Metrics

The following metrics are available for each SQL Server instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
CPU Percentage	Metric	Percent	Average	The average percentage of CPU used in the SQL Server databases.
SQL Version	Property	String	Not applicable.	The version of the SQL Server.
Data IO Percentage	Metric	Percent	Average	The average percentage of data IO used in the SQL Server databases.
DTU Used	Metric	Count	Average	The average number of DTUs used in the DTU-based SQL Server databases.

Name	Type	Unit	Aggregation Type	Description
In-Memory OLTP Storage Percent	Metric	Percent	Average	The average percentage of in-memory OLTP storage in the SQL Server databases.
Log IO Percentage	Metric	Percent	Average	The average percentage of log IO used in the SQL Server databases.
Sessions Percentage	Metric	Percent	Average	The average percentage of sessions in the SQL Server databases.
Workers Percentage	Metric	Percent	Average	The average percentage of workers in the SQL Server databases.

SQL Database Metrics

The following metrics are available for each SQL Database instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
CPU Percentage	Metric	Percent	Average	The percentage of CPU in use.
Data IO Percentage	Metric	Percent	Average	The percentage of data IO in use.
Log IO Percentage	Metric	Percent	Average	The percentage of log IO in use. Not applicable to data warehouses.
DTU Percentage	Metric	Percent	Average	The percentage of DTU in use. Applies to DTU-based databases.
Data Space Used	Metric	Bytes	Maximum	The total size of the database. Not applicable to data warehouses.
Successful Connections	Metric	Count	Total	The number of successful connections to the database.

Name	Type	Unit	Aggregation Type	Description
Failed Connections	Metric	Count	Total	The number of failed connections to the database.
Blocked by Firewall	Metric	Count	Total	The number of connections to the database blocked by firewall.
Deadlocks	Metric	Count	Total	The number of deadlocks. Not applicable to data warehouses.
Data Space Used Percent	Metric	Percent	Maximum	The percentage of database size. Not applicable to data warehouses or hyper-scale databases.
In-Memory OLTP Storage Percent	Metric	Percent	Average	The percentage of in-memory OLTP storage. Not applicable to data warehouses.
Workers Percentage	Metric	Percent	Average	The percentage of workers. Not applicable to data warehouses.
Sessions Percentage	Metric	Percent	Average	The percentage of sessions. Not applicable to data warehouses.
DTU Limit	Metric	Count	Average	The maximum number of DTUs. Applies to DTU-based databases.
DTU Used	Metric	Count	Average	The number of DTUs used. Applies to DTU-based databases.
CPU Limit	Metric	Count	Average	The maximum number of CPUs. Applies to vCore-based databases.
CPU Used	Metric	Count	Average	The number of CPUs used. Applies to vCore-based databases.

Name	Type	Unit	Aggregation Type	Description
DWU Limit	Metric	Count	Maximum	The maximum number of DWUs. Applies only to data warehouses.
DWU Percentage	Metric	Percent	Maximum	The percentage of DWUs used. Applies only to data warehouses.
DWU Used	Metric	Count	Maximum	The number of DWUs used. Applies only to data warehouses.
DW Node Level CPU Percentage	Metric	Percent	Average	The DW node level CPU percentage.
DW Node Level Data IO Percentage	Metric	Percent	Average	The DW node level Data IO percentage.
Cache Hit Percentage	Metric	Percent	Maximum	The percentage of cache hits. Applies only to data warehouses.
Cache Used Percentage	Metric	Percent	Maximum	The percentage of cache used. Applies only to data warehouses.
Local tempdb Percentage	Metric	Percent	Average	The local <i>tempdb</i> percentage. Applies only to data warehouses.
App CPU Billed	Metric	Count	Total	The number of app CPUs billed. Applies to server-less databases.
App CPU Percentage	Metric	Percent	Average	The app CPU percentage. Applies to server-less databases.
App Memory Used Percentage	Metric	Percent	Average	The percentage of app memory used. Applies to server-less databases.
Data Space Allocated	Metric	Bytes	Average	The data space allocated. Not applicable to data warehouses.

MySQL Server Metrics

The following metrics are available for each MySQL Server instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
CPU Percent	Metric	Percent	Average	The percentage of CPU in use.
Memory Percent	Metric	Percent	Average	The percentage of memory in use.
IO Percent	Metric	Percent	Average	The percentage of IO in use.
Storage Percent	Metric	Percent	Average	The percentage of storage used out of the server's maximum.
Storage Used	Metric	Bytes	Average	The amount of storage in use. The storage used by the service includes the database files, transaction logs, and the server logs.
Storage Limit	Metric	Bytes	Average	The maximum storage for the server.
Server Log Storage Percent	Metric	Percent	Average	The percentage of server log storage used out of the server's maximum server log storage.
Server Log Storage Used	Metric	Bytes	Average	The amount of server log storage in use.
Server Log Storage Limit	Metric	Bytes	Average	The maximum server log storage for the server.
Active Connections	Metric	Count	Average	The number of active connections to the server.
Failed Connections	Metric	Count	Total	The number of failed connections to the server.

Name	Type	Unit	Aggregation Type	Description
Replication Lag in Seconds	Metric	Seconds	Average	The number of seconds the replica server is lagging against the primary server.
Backup Storage Used	Metric	Bytes	Average	The amount of backup storage used.
Network Out	Metric	Bytes	Total	The Network Out across active connections.
Network In	Metric	Bytes	Total	The Network In across active connections.

PostgreSQL Server Metrics

The following metrics are available for each PostgreSQL Server instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
CPU Percent	Metric	Percent	Average	The percentage of CPU in use.
Memory Percent	Metric	Percent	Average	The percentage of memory in use.
IO Percent	Metric	Percent	Average	The percentage of IO in use.
Storage Percent	Metric	Percent	Average	The percentage of storage used out of the server's maximum.
Storage Used	Metric	Bytes	Average	The amount of storage in use. The storage used by the service includes the database files, transaction logs, and the server logs.
Storage Limit	Metric	Bytes	Average	The maximum storage for the server.

Name	Type	Unit	Aggregation Type	Description
Server Log Storage Percent	Metric	Percent	Average	The percentage of server log storage used out of the server's maximum server log storage.
Server Log Storage Used	Metric	Bytes	Average	The amount of server log storage in use.
Server Log Storage Limit	Metric	Bytes	Average	The maximum server log storage for the server.
Active Connections	Metric	Count	Average	The number of active connections to the server.
Failed Connections	Metric	Count	Total	The number of failed connections to the server.
Backup Storage Used	Metric	Bytes	Average	The amount of backup storage used.
Network Out	Metric	Bytes	Total	The Network Out across active connections.
Network In	Metric	Bytes	Total	The Network In across active connections.
Replica Lag	Metric	Seconds	Maximum	The number of seconds the replica server is lagging against the primary server.
Max Lag Across Replicas	Metric	Bytes	Maximum	The lag in bytes of the most lagging replica server.

Network Interface Metrics

The following metrics are available for each Network Interface instance of the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
Bytes Sent	Metric	Count	Total	The number of bytes the Network Interface sent.
Bytes Received	Metric	Count	Total	The number of bytes the Network Interface received.
Packets Sent	Metric	Count	Total	The number of packets the Network Interface sent.
Packets Received	Metric	Count	Total	The number of packets the Network Interface received.

Load Balancer Metrics

The following metrics are available for each Load Balancer instance for the Management Pack for Microsoft Azure in vRealize Operations Manager .

For more information about each metric, see the Microsoft Azure documentation at <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/metrics-supported>.

Name	Type	Unit	Aggregation Type	Description
Data Path Availability	Metric	Count	Average	The average Load Balancer data path availability per time duration.
Health Probe Status	Metric	Count	Average	The average Load Balancer health probe status per time duration.
Byte Count	Metric	Count	Total	The total number of bytes transmitted within a time period.
Packet Count	Metric	Count	Total	The total number of packets transmitted within a time period.

Metrics for Management Pack for AWS

The Management Pack for AWS imports Amazon ElastiCache metrics which collect data for vRealize Operations Manager components.

EC2 Metrics

The following metrics are available for each EC2 instance in your vRealize Operations Manager environment.

Note Capacity calculations are enabled by the default policy and these calculations are based on the CPU and Memory utilization metrics.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/ec2-metricscollected.html>.

Table 8-132. EC2 Metrics

Name	Category	Type	Unit	Instanced
DiskReadOps	Disk Space	Metric	Count	No
DiskWriteOps	Disk Space	Metric	Count	No
DiskReadBytes	Disk Space	Metric	Bytes	No
DiskWriteBytes	Disk Space	Metric	Bytes	No
Disk I/O	Disk Space	Metric	Count	No
CPUUtilization	CPU	Metric	Percent	No
CPUCreditUsage	CPU	Metric	Count	No
CPUCreditBalance	CPU	Metric	Count	No
NetworkIn	Network	Metric	Bytes	No
NetworkOut	Network	Metric	Bytes	No
NetworkPacketsIn	Network	Metric	Count	No
NetworkPacketsOut	Network	Metric	Count	No
Network I/O	Network	Metric	Count	No
StatusCheckFailed	Status	Metric	Count	No
StatusCheckFailed_Instance	Status	Metric	Count	No
StatusCheckFailed_System	Status	Metric	Count	No
Runtime	Status	Metric	Hours	No
Memory Available	Memory	Metric	Megabytes	No
MemoryUsed	Memory	Metric	Megabytes	No
MemoryUtilization	Memory	Metric	Percent	No
SwapUsed	Memory	Metric	Megabytes	No

Table 8-132. EC2 Metrics (continued)

Name	Category	Type	Unit	Instanced
SwapUtilization	Memory	Metric	Percent	No
pagefileAvailable	Memory	Metric	Megabytes	No
pagefileUsed	Memory	Metric	Megabytes	No
pagefileUtilization	Memory	Metric	Percent	No
DiskSpaceAvailable	Filesystem	Metric	Gigabytes	No
DiskSpaceUsed	Filesystem	Metric	Gigabytes	No
DiskSpaceUtilization	Filesystem	Metric	Percent	No
VolumAvailable	Filesystem	Metric	Gigabytes	No
VolumeUsed	Filesystem	Metric	Gigabytes	No
VolumeUtilization	Filesystem	Metric	Percent	No
sec	Perfmon	Metric	Count	No
Processor Queue Length	Perfmon	Metric	Count	No

EC2 Volume Metrics

The following metrics are available for each EC2 Volume instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

Table 8-133. EC2 Volume Metrics

Name	Category	Type	Unit	Instanced
VolumeReadBytes	Disk Space	Metric	Bytes	No
VolumeWriteBytes	Disk Space	Metric	Bytes	No
VolumeReadOps	Disk Space	Metric	Count	No
VolumeWriteOps	Disk Space	Metric	Count	No
VolumeTotalReadTime	Disk Space	Metric	Seconds	No
VolumeTotalWriteTime	Disk Space	Metric	Seconds	No
VolumeIdleTime	Disk Space	Metric	Seconds	No
VolumeQueueLength	Disk Space	Metric	Count	No

Table 8-133. EC2 Volume Metrics (continued)

Name	Category	Type	Unit	Instanced
VolumeThroughputPercentage	Disk Space	Metric	Percent	No
VolumeConsumedReadWriteOps	Disk Space	Metric	Count	No
VolumeCapacity	Disk Space	Metric	Count	No

EC2 Load Balancer Metrics

The following metrics are available for each EC2 Load Balancer instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/US_MonitoringLoadBalancerWithCW.html

Table 8-134. EC2 Load Balancer Metrics

Name	Category	Type	Unit	Instanced
Latency	General	Metric	Seconds	No
RequestCount	General	Metric	Count	No
HealthyHostCount	General	Metric	Count	No
UnHealthyHostCount	General	Metric	Count	No
HTTPCode_ELB_4XX	General	Metric	Count	No
HTTPCode_ELB_5XX	General	Metric	Count	No
HTTPCode_Backend_2XX	General	Metric	Count	No
HTTPCode_Backend_3XX	General	Metric	Count	No
HTTPCode_Backend_4XX	General	Metric	Count	No
HTTPCode_Backend_5XX	General	Metric	Count	No
BackendConnectionErrors	General	Metric	Count	No
SurgeQueueLength	General	Metric	Count	No
SpilloverCount	General	Metric	Count	No

Network Load Balancer Metrics

The following metrics are available for each Network Load Balancer instance in your vRealize Operations Manager environment.

Table 8-135. Network Load Balancer Metrics

Name	Category	Type	Unit	Instanced
HealthyHostCount	General	Metric	Count	No
UnHealthyHostCount	General	Metric	Count	No
ActiveFlowCount	General	Metric	Count	No
ConsumedLCUs	General	Metric	Count	No
NewFlowCount	General	Metric	Count	No
ProcessedBytes	General	Metric	Bytes	No
TCP_Client_Reset_Count	General	Metric	Count	No
TCP_ELB_Reset_Count	General	Metric	Count	No
TCP_Target_Reset_Count	General	Metric	Count	No

Application Load Balancer Metrics

The following metrics are available for each Application Load Balancer instance in your vRealize Operations Manager environment.

Table 8-136. Application Load Balancer Metrics

Name	Category	Type	Unit	Instanced
ActiveConnectionCount	General	Metric	Count	No
ConsumedLCUs	General	Metric	Count	No
ClientTLSNegotiationErrorCount	General	Metric	Count	No
Latency	General	Metric	Seconds	No
RequestCount	General	Metric	Count	No
HealthyHostCount	General	Metric	Count	No
UnHealthyHostCount	General	Metric	Count	No
HTTPCode_ELB_4XX_Count	General	Metric	Count	No

Table 8-136. Application Load Balancer Metrics (continued)

Name	Category	Type	Unit	Instanced
HTTPCode_ELB_5XX_Count	General	Metric	Count	No
HTTPCode_Target_2XX_Count	General	Metric	Count	No
HTTPCode_Target_3XX_Count	General	Metric	Count	No
HTTPCode_Target_4XX_Count	General	Metric	Count	No
HTTPCode_Target_5XX_Count	General	Metric	Count	No
IPv6ProcessedBytes	General	Metric	Bytes	No
IPv6RequestCount	General	Metric	Count	No
NewConnectionCount	General	Metric	Count	No
RejectedConnectionCount	General	Metric	Count	No
ProcessedBytes	General	Metric	Bytes	No
RuleEvaluations	General	Metric	Count	No
TargetResponseTime	General	Metric	Seconds	No
TargetTLSNegotiationErrorCount	General	Metric	Count	No

EC2 Auto Scale Group Metrics

The following metrics are available for each EC2 Auto Scale Group instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-instance-monitoring.html>

Table 8-137. EC2 Auto Scale Group Metrics

Name	Category	Type	Unit	Instanced
GroupMinSize	General	Metric	Count	No
GroupMaxSize	General	Metric	Count	No
GroupDesiredCapacity	General	Metric	Count	No
GroupInServiceInstances	General	Metric	Count	No

Table 8-137. EC2 Auto Scale Group Metrics (continued)

Name	Category	Type	Unit	Instanced
GroupPendingInstances	General	Metric	Count	No
GroupTerminatingInstances	General	Metric	Count	No
GroupTotalInstances	General	Metric	Count	No
DiskReadOps	Disk	Metric	Count	No
DiskWriteOps	Disk	Metric	Count	No
DiskReadBytes	Disk	Metric	Bytes	No
DiskWriteBytes	Disk	Metric	Bytes	No
Aggregate Disk I/O	Disk	Metric	Bytes	No
Aggregate Disk I/O	Disk	Metric	Count	No
CPUUtilization	CPU	Metric	Percent	No
NetworkIn	Network	Metric	Bytes	No
NetworkOut	Network	Metric	Bytes	No
StatusCheckFailed	Status	Metric	Count	No
StatusCheckFailed_Instance	Status	Metric	Count	No
StatusCheckFailed_System	Status	Metric	Count	No

EMR Job Flow Metrics

The following metrics are available for each EMR Job Flow instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/emr-metricscollected.html>

Table 8-138. EMR Job Flow Metrics

Name	Category	Type	Unit	Instanced
CoreNodesPending	Health	Metric	Count	No
CoreNodesRunning	Health	Metric	Count	No
JobsFailed	Health	Metric	Count	No
JobsRunning	Health	Metric	Count	No
LiveDataNodes	Health	Metric	Percent	No

Table 8-138. EMR Job Flow Metrics (continued)

Name	Category	Type	Unit	Instanced
LiveTaskTrackers	Health	Metric	Percent	No
MissingBlocks	Health	Metric	Count	No
TaskNodesPending	Health	Metric	Count	No
TaskNodesRunning	Health	Metric	Count	No
TotalLoad	Health	Metric	Count	No
CapacityRemaining GB	Health	Metric	Count	No
CorruptBlocks	Health	Metric	Count	No
PendingDeletionBlocks	Health	Metric	Count	No
UnderReplicatedBlocks	Health	Metric	Count	No
dfs.FSNamesystem. PendingReplication Blocks	Health	Metric	Count	No
HDFSBytesRead	Performance and Progress	Metric	Count	No
HDFSBytesWritten	Performance and Progress	Metric	Count	No
HDFSUtilization	Performance and Progress	Metric	Percent	No
ISIdle	Performance and Progress	Metric	Count	No
MapSlotsOpen	Performance and Progress	Metric	Percent	No
ReduceSlotsOpen	Performance and Progress	Metric	Percent	No
RemainingMapTasks	Performance and Progress	Metric	Count	No
RemainingMapTasks PerSlot	Performance and Progress	Metric	Ratio	No
RemainingReduceTasks	Performance and Progress	Metric	Count	No
RunningMapTasks	Performance and Progress	Metric	Count	No
RunningReduceTasks	Performance and Progress	Metric	Count	No
S3BytesRead	Performance and Progress	Metric	Count	No

Table 8-138. EMR Job Flow Metrics (continued)

Name	Category	Type	Unit	Instanced
S3BytesWritten	Performance and Progress	Metric	Count	No
HBaseMostRecentBackupDuration	HBase Backups	Metric	Minutes	No
HBaseTimeSinceLastSuccessfulBackup	HBase Backups	Metric	Minutes	No

Entity Status Metrics

The following metrics are available for each Entity Status instance in your vRealize Operations Manager environment.

Table 8-139. Entity Status Metrics

Name	Category	Type	Unit	Instanced
Total EC2 Instances	General	Metric		No
Active EC2 Instances	General	Metric		No
Number of S3 Buckets	General	Metric		No
Number of EC2 Volumes	General	Metric		No
Number of Load Balancers	General	Metric		No
Number of Auto Scaling Groups	General	Metric		No
Number of EMR Job Flows	General	Metric		No
Number of ElastiCache Clusters	General	Metric		No
Number of ElastiCache Nodes	General	Metric		No
Number of RDS DB Instances	General	Metric		No
Number of Lambda Functions	General	Metric		No
Number of Redshift Clusters	General	Metric		No
Number of Redshift Nodes	General	Metric		No
Number of ECR Repositories	General	Metric		No

Table 8-139. Entity Status Metrics (continued)

Name	Category	Type	Unit	Instanced
Number of ECR Images	General	Metric		No
Number of SQS Queues	General	Metric		No
Number of WorkSpaces	General	Metric		No
Number of ECS Clusters	General	Metric		No
Number of ECS Services	General	Metric		No
Number of DynamoDB Tables	General	Metric		No
Number of DynamoDB Accelerator Clusters	General	Metric		No
Number of DynamoDB Accelerator Nodes	General	Metric		No
Number of VPC NAT Gateways	General	Metric		No
Number of Application Load Balancers	General	Metric		No
Number of CloudFormation Stacks	General	Metric		No
Number of Network Load Balancers	General	Metric		No
Number of Classic Load Balancers	General	Metric		No
Number of Security Groups	General	Metric		No
Number of Elastic IPs	General	Metric		No
Number of CloudFront Distribution	General	Metric		No

ElastiCache Cache Node Metrics

The following metrics are available for each ElastiCache Cache Node instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.Redis.html>, <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.HostLevel.html>, and <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.Memcached.html>.

Table 8-140. ElastiCache Cache Node Metrics

Name	Category	Type	Unit	Instanced
CPUUtilization	CPU	Metric	Percent	No
SwapUsage	Memory	Metric	Bytes	No
FreeableMemory	Memory	Metric	Bytes	No
NetworkBytesIn	Network	Metric	Bytes	No
NetworkBytesOut	Network	Metric	Bytes	No
BytesUsedForCacheItems	Memory	Metric	Bytes	No
BytesReadIntoMemcached	Memory	Metric	Bytes	No
BytesWrittenOutFromMemM	Memory	Metric	Bytes	No
BytesUsedForHash	Memory	Metric	Bytes	No
BytesUsedForCache	Memory	Metric	Bytes	No
CasBadval	Memory	Metric	Count	No
CasHits	Memory	Metric	Count	No
CasMisses	Memory	Metric	Count	No
UnusedMemory	Memory	Metric	Count	No
CmdFlush	Commands	Metric	Count	No
CmdGet	Commands	Metric	Count	No
CmdSet	Commands	Metric	Count	No
CmdConfigGet	Commands	Metric	Count	No
CmdConfigSet	Commands	Metric	Count	No
CmdTouch	Commands	Metric	Count	No
GetTypeCmds	Commands	Metric	Count	No
SetTypeCmds	Commands	Metric	Count	No
KeyBasedCmds	Commands	Metric	Count	No

Table 8-140. ElastiCache Cache Node Metrics (continued)

Name	Category	Type	Unit	Instanced
StringBasedCmds	Commands	Metric	Count	No
HashBasedCmds	Commands	Metric	Count	No
ListBasedCmds	Commands	Metric	Count	No
SetBasedCmds	Commands	Metric	Count	No
SortedSetBasedCmds	Commands	Metric	Count	No
CurrConnections	Performance	Metric	Count	No
CurrItems	Performance	Metric	Count	No
DecrHits	Performance	Metric	Count	No
DecrMisses	Performance	Metric	Count	No
DeleteHits	Performance	Metric	Count	No
DeleteMisses	Performance	Metric	Count	No
Evictions	Performance	Metric	Count	No
GetHits	Performance	Metric	Count	No
GetMisses	Performance	Metric	Count	No
IncrHits	Performance	Metric	Count	No
IncrMisses	Performance	Metric	Count	No
Reclaimed	Performance	Metric	Count	No
CurrConfig	Performance	Metric	Count	No
EvictedUnfetched	Performance	Metric	Count	No
ExpiredUnfetched	Performance	Metric	Count	No
SlabsMoved	Performance	Metric	Count	No
TouchHits	Performance	Metric	Count	No
TouchMisses	Performance	Metric	Count	No
NewConnections	Performance	Metric	Count	No
NewItems	Performance	Metric	Count	No
CacheHits	Performance	Metric	Count	No

Table 8-140. ElastiCache Cache Node Metrics (continued)

Name	Category	Type	Unit	Instanced
CacheMisses	Performance	Metric	Count	No
ReplicationLag	Performance	Metric	Count	No

RDS DB Instance Metrics

The following metrics are available for each RDS DB instance in your vRealize Operations Manager environment.

Table 8-141. RDS DB Instance Metrics

Name	Category	Type	Unit	Instanced
CPUUtilization	CPU	Metric	Percent	No
CPUCreditUsage	CPU	Metric	Count	No
CPUCreditBalance	CPU	Metric	Count	No
FreeableMemory	Memory	Metric	Bytes	No
BinLogDiskUsage	Disk	Metric	Bytes	No
DiskQueueDepth	Disk	Metric	Count	No
FreeStorageSpace	Disk	Metric	Bytes	No
SwapUsage	Disk	Metric	Bytes	No
ReadIOPS	Disk	Metric	Count/second	No
WriteIOPS	Disk	Metric	Count/second	No
ReadLatency	Disk	Metric	Seconds	No
WriteLatency	Disk	Metric	Seconds	No
ReadThroughput	Disk	Metric	Bytes/seconds	No
WriteThroughput	Disk	Metric	Bytes/seconds	No
DatabaseConnections	Performance	Metric	Count	No

Lambda Metrics

The following metrics are available for each Lambda instance in your vRealize Operations Manager environment.

Table 8-142. Lamda Metrics

Name	Category	Type	Unit	Instanced
Invocations	General	Metric	Count	No
Errors	General	Metric	Count	No
Duration	General	Metric	Milliseconds	No
Throttles	General	Metric	Count	No
IteratorAge	General	Metric	Milliseconds	No

Redshift Cluster Metrics

The following metrics are available for each Redshift Cluster instance in your vRealize Operations Manager environment.

Table 8-143. Redshift Cluster Metrics

Name	Category	Type	Unit	Instanced
CPUUtilization Average	CPU	Metric	Percent	No
DatabaseConnections	General	Metric	Count	No
HealthStatus	General	Metric	Count	No
MaintenanceMode	General	Metric	Count	No
PercentageDiskSpaceUsed	Disk	Metric	Percent	No
ReadIOPS	Disk	Metric	Count/second	No
ReadLatency	Disk	Metric	Count/second	No
ReadThroughput	Disk	Metric	Bytes/second	No
WriteIOPS	Disk	Metric	Count/second	No
WriteLatency	Disk	Metric	Seconds	No
WriteThroughput	Disk	Metric	Bytes/second	No
NetworkReceiveThroughput	Network	Metric	Bytes/second	No
NetworkTransmitThroughput	Network	Metric	Bytes/second	No

Redshift Node Metrics

The following metrics are available for each Redshift Node instance in your vRealize Operations Manager environment.

Table 8-144. Redshift Node Metrics

Name	Category	Type	Unit	Instanced
CPUUtilizationAverage	CPU	Metric	Percent	No
DatabaseConnections	General	Metric	Count	No
HealthStatus	General	Metric	Count	No
MaintenanceMode	General	Metric	Count	No
PercentageDiskSpaceUsed	Disk	Metric	Percent	No
ReadIOPS	Disk	Metric	Count/second	No
ReadLatency	Disk	Metric	Count/second	No
ReadThroughput	Disk	Metric	Bytes/second	No
WriteIOPS	Disk	Metric	Count/second	No
WriteLatency	Disk	Metric	Seconds	No
WriteThroughput	Disk	Metric	Bytes/second	No
NetworkReceiveThroughput	Network	Metric	Bytes/second	No
NetworkTransmitThroughput	Network	Metric	Bytes/second	No

AWS Workspace Metrics

The following metrics are available for each AWS Workspace instance in your vRealize Operations Manager environment.

Table 8-145. AWS Workspace Metrics

Name	Category	Type	Unit	Instanced
Available	General	Metric	Count	No
Unhealthy	General	Metric	Count	No
ConnectionAttempt	General	Metric	Count	No
ConnectionSuccess	General	Metric	Count	No
ConnectionFailure	General	Metric	Count	No
SessionDisconnect	General	Metric	Count	No
UserConnected	General	Metric	Count	No
Stopped	General	Metric	Count	No

Table 8-145. AWS Workspace Metrics (continued)

Name	Category	Type	Unit	Instanced
Maintenance	General	Metric	Count	No
SessionLaunchTime	General	Metric	Seconds	No
InSessionLatency	General	Metric	Milliseconds	No

ECS Cluster Metrics

The following metrics are available for each ECS Cluster instance in your vRealize Operations Manager environment.

Table 8-146. ECS Cluster Metrics

Name	Category	Type	Unit	Instanced
CPUReservation Average	CPU	Metric	Percent	No
CPUUtilization	CPU	Metric	Percent	No
MemoryReservation	Memory	Metric	Percent	No
MemoryUtilization	Memory	Metric	Percent	No

ECS Service Metrics

The following metrics are available for each ECS Service instance in your vRealize Operations Manager environment.

Table 8-147. ECS Service Metrics

Name	Category	Type	Unit	Instanced
CPUReservation Average	CPU	Metric	Percent	No
CPUUtilization	CPU	Metric	Percent	No
MemoryReservation	Memory	Metric	Percent	No
MemoryUtilization	Memory	Metric	Percent	No

DynamoDB Metrics

The following metrics are available for each DynamoDB instance in your vRealize Operations Manager environment.

Table 8-148. DynamoDB Metrics

Name	Category	Type	Unit	Instanced
ConditionalCheckFailedRequests	General	Metric	Count	No
ConsumedReadCapacityUnits	General	Metric	Count	No
ConsumedWriteCapacityUnits	General	Metric	Count	No
OnlineIndexConsumedWriteCapacity	General	Metric	Count	No
OnlineIndexPercentageProgress	General	Metric	Count	No
OnlineIndexThrottleEvents Average	General	Metric	Count	No
ReadThrottleEvents	General	Metric	Count	No
ReturnedBytes Average	General	Metric	Count	No
ReturnedItemCount	General	Metric	Count	No
ReturnedRecordsCount	General	Metric	Count	No
SuccessfulRequestLatency	General	Metric	Count	No
SystemErrors	General	Metric	Count	No
TimeToLiveDeletedItemCount	General	Metric	Count	No
ThrottledRequests	General	Metric	Count	No
UserErrors	General	Metric	Count	No
WriteThrottleEvents Average	General	Metric	Count	No
ProvisionedReadCapacityUnits	General	Metric	Count	No
ProvisionedWriteCapacityUnit	General	Metric	Count	No

S3 Bucket Metrics

The following metrics are available for each S3 Bucket instance in your vRealize Operations Manager environment.

Table 8-149. S3 Bucket Metrics

Name	Category	Type	Unit	Instanced
BucketSizeBytes Average	General	Metric	Bytes	No
BucketSizeBytes Average	General	Metric	Count	No
AllRequests Average	General	Metric	Count	No
GetRequests Average	General	Metric	Count	No
PutRequests Average	General	Metric	Count	No
DeleteRequests Average	General	Metric	Count	No
HeadRequests Average	General	Metric	Count	No
PostRequests Average	General	Metric	Count	No
ListRequests Average	General	Metric	Count	No
BytesDownloaded Average	General	Metric	Bytes	No
BytesUploaded Average	General	Metric	Bytes	No
4xxErrors	General	Metric	Count	No
5xxErrors	General	Metric	Count	No
FirstByteLatency	General	Metric	Milliseconds	No
TotalRequestLatency	General	Metric	Milliseconds	No

VPC Nat Gateway Metrics

The following metrics are available for each VPC Nat Gateway instance in your vRealize Operations Manager environment.

Table 8-150. VPC Nat Gateway Metrics

Name	Category	Type	Unit	Instanced
ErrorPortAllocation	General	Metric	Count	No
ActiveConnectionCount	General	Metric	Count	No
ConnectionAttemptCount	General	Metric	Count	No

Table 8-150. VPC Nat Gateway Metrics (continued)

Name	Category	Type	Unit	Instanced
ConnectionEstablishedCount	General	Metric	Count	No
IdleTimeoutCount	General	Metric	Count	No
PacketsOutToDestination	Network	Metric	Count	No
PacketsOutToSource	Network	Metric	Count	No
PacketsInFromSource	Network	Metric	Count	No
PacketsInFromDestination	Network	Metric	Count	No
BytesOutToDestination	Network	Metric	Bytes	No
BytesOutToSource	Network	Metric	Bytes	No
BytesInFromSource	Network	Metric	Bytes	No
BytesInFromDestination	Network	Metric	Bytes	No
PacketsDropCount	Network	Metric	Count	No

Dax Cluster Metrics

The following metrics are available for each Dax Cluster instance in your vRealize Operations Manager environment.

Table 8-151. DAX Cluster Metrics

Name	Category	Type	Unit	Instanced
ItemCacheMisses	General	Metric	Count	No
QueryCacheHits	General	Metric	Count	No
ScanCacheHits	General	Metric	Count	No
FailedRequestCount	General	Metric	Count	No
ScanCacheMisses	General	Metric	Count	No
ErrorRequestCount	General	Metric	Count	No
QueryCacheMisses	General	Metric	Count	No
TotalRequestCount	General	Metric	Count	No
EstimatedDbSize	General	Metric	Bytes	No

Table 8-151. DAX Cluster Metrics (continued)

Name	Category	Type	Unit	Instanced
EvictedSize	General	Metric	Bytes	No
FaultRequestCount	General	Metric	Count	No
ScanRequestCount	General	Metric	Count	No
ItemCacheHits	General	Metric	Count	No
QueryRequestCount	General	Metric	Count	No
DeleteItemRequestCount	General	Metric	Count	No
GetItemRequestCount	General	Metric	Count	No
UpdateItemRequestCount	General	Metric	Count	No
BatchWriteItemRequestCount	General	Metric	Count	No
PutItemRequestCount	General	Metric	Count	No
BatchGetItemRequestCount	General	Metric	Count	No
PutItemRequestCount	General	Metric	Count	No

DAX Node Metrics

The following metrics are available for each DAX node instance in your vRealize Operations Manager environment.

Table 8-152. DAX Node Metrics

Name	Category	Type	Unit	Instanced
ItemCacheMisses	General	Metric	Count	No
QueryCacheHits	General	Metric	Count	No
ScanCacheHits	General	Metric	Count	No
FailedRequestCount	General	Metric	Count	No
ScanCacheMisses	General	Metric	Count	No
ErrorRequestCount	General	Metric	Count	No
QueryCacheMisses	General	Metric	Count	No
TotalRequestCount	General	Metric	Count	No

Table 8-152. DAX Node Metrics (continued)

Name	Category	Type	Unit	Instanced
EstimatedDbSize	General	Metric	Bytes	No
EvictedSize	General	Metric	Bytes	No
FaultRequestCount	General	Metric	Count	No
ScanRequestCount	General	Metric	Count	No
ItemCacheHits	General	Metric	Count	No
QueryRequestCount	General	Metric	Count	No
DeleteItemRequestCount	General	Metric	Count	No
GetItemRequestCount	General	Metric	Count	No
UpdateItemRequestCount	General	Metric	Count	No
BatchWriteItemRequestCount	General	Metric	Count	No
PutItemRequestCount	General	Metric	Count	No
BatchGetItemRequestCount	General	Metric	Count	No
PutItemRequestCount	General	Metric	Count	No

Direct Connect Metrics

The following metrics are available for each Direct Connect instance in your vRealize Operations Manager environment.

Table 8-153. Direct Connect Metrics

Name	Category	Type	Unit	Instanced
ConnectionState	General	Metric	Count	No
ConnectionBpsEgress	General	Metric	Bits/Second	No
ConnectionBpsIngress	General	Metric	Bits/Second	No
ConnectionPpsEgress	General	Metric	Count/Second	No
ConnectionPpsIngress	General	Metric	Count/Second	No

Table 8-153. Direct Connect Metrics (continued)

Name	Category	Type	Unit	Instanced
ConnectionCRCErrorCount	General	Metric	Count	No
ConnectionLightLevelTx	General	Metric	dBm	No
ConnectionLightLevelRx	General	Metric	dBm	No

Health Check Metrics

The following metrics are available for each Health Check instance in your vRealize Operations Manager environment.

Table 8-154. Health Check Metrics

Name	Category	Type	Unit	Instanced
ChildHealthCheckHealthyCount		Metric	Count	No
ConnectionTime		Metric	Milliseconds	No
HealthCheckPercentageHealthy		Metric	Percent	No
SSLHandshakeTime		Metric	Milliseconds	No
TimeToFirstByte		Metric	Milliseconds	No

ElastiCache Cache Cluster Metrics

The following metrics are available for each ElastiCache Cache Cluster instance in your vRealize Operations Manager environment.

For a description of each metric, see the Amazon Web Service documentation at <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.Redis.html> and <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheMetrics.HostLevel.html>.

Table 8-155. ElasticCache Cache Cluster Metrics

Name	Category	Type	Unit	Instanced
CPUUtilization	CPU	Metric	Percent	No
NetworkBytesIn	Network	Metric	Bytes	No
NetworkBytesOut	Network	Metric	Bytes	No
SwapUsage	Memory	Metric	Bytes	No
FreeableMemory	Memory	Metric	Bytes	No
BytesUsedForCache	Memory	Metric	Bytes	No

Table 8-155. ElasticCache Cache Cluster Metrics (continued)

Name	Category	Type	Unit	Instanced
GetTypeCmds	Commands	Metric	Count	No
SetTypeCmds	Commands	Metric	Count	No
KeyBasedCmds	Commands	Metric	Count	No
StringBasedCmds	Commands	Metric	Count	No
HashBasedCmds	Commands	Metric	Count	No
ListBasedCmds	Commands	Metric	Count	No
SetBasedCmds	Commands	Metric	Count	No
SortedSetBasedCmds	Commands	Metric	Count	No
CurrConnections	Performance	Metric	Count	No
CurrItems	Performance	Metric	Count	No
Evictions	Performance	Metric	Count	No
Reclaimed	Performance	Metric	Count	No
NewConnections	Performance	Metric	Count	No
NewItems	Performance	Metric	Count	No
CacheHits	Performance	Metric	Count	No
CacheMisses	Performance	Metric	Count	No
ReplicationLag	Performance	Metric	Count	No

EFS Metrics

The following metrics are available for each EFS instance in your vRealize Operations Manager environment.

Table 8-156. EFS Metrics

Service	Metrics
EFS	BurstCreditBalance
	ClientConnections
	DataReadIOBytes
	DataWriteIOBytes
	MetadataIOBytes
	PercentIOLimit

Table 8-156. EFS Metrics (continued)

Service	Metrics
	PermittedThroughput
	TotalIOBytes

Elastic Beanstalk Environment Metrics

The following metrics are available for each Elastic Beanstalk Environment instance in your vRealize Operations Manager environment.

Table 8-157. Elastic Beanstalk Environment Metrics

Service	Metrics
Elastic Beanstalk Environment	InstancesSevere
	InstancesDegraded
	ApplicationRequests5xx
	ApplicationRequests4xx
	ApplicationLatencyP50
	ApplicationLatencyP95
	ApplicationLatencyP85
	InstancesUnknown
	ApplicationLatencyP90
	InstancesInfo
	InstancesPending
	ApplicationLatencyP75
	ApplicationLatencyP10
	ApplicationLatencyP99
	ApplicationRequestsTotal
	InstancesNoData
	ApplicationLatencyP99.9
	ApplicationRequests3xx
	ApplicationRequests2xx
	InstancesOk

Table 8-157. Elastic Beanstalk Environment Metrics (continued)

Service	Metrics
	InstancesWarning
	EnvironmentHealth

AWS Transit Gateway Metrics

The following metrics are available for each AWS Transit Gateway instance in your vRealize Operations Manager environment.

Table 8-158. AWS Transit Gateway Metrics

Service	Metrics
AWS Transit Gateway	BytesIn
	BytesOut
	PacketsIn
	PacketsOut
	PacketDropCountBlackhole
	PacketDropCountNoRoute
	BytesDropCountNoRoute
	BytesDropCountBlackhole

EKS Cluster Metrics

The following metrics are available for each EKS Cluster instance in your vRealize Operations Manager environment.

Table 8-159. EKS Cluster Metrics

Service	Metrics
EKS Cluster	cluster_failed_node_count
	cluster_node_count
	namespace_number_of_running_pods
	node_cpu_limit
	node_cpu_reserved_capacity
	node_cpu_usage_total
	node_cpu_utilization
	node_filesystem_utilization

Table 8-159. EKS Cluster Metrics (continued)

Service	Metrics
	node_memory_limit
	node_memory_reserved_capacity
	node_memory_utilization
	node_memory_working_set
	node_network_total_bytes
	node_number_of_running_containers
	node_number_of_running_pods
	pod_cpu_reserved_capacity
	pod_cpu_utilization
	pod_cpu_utilization_over_pod_limit
	pod_memory_reserved_capacity
	pod_memory_utilization
	pod_memory_utilization_over_pod_limit
	pod_number_of_container_restarts
	pod_network_rx_bytes
	pod_network_tx_bytes
	service_number_of_running_pods

Metrics in VMware Cloud on AWS

The VMware Cloud on AWS collects metrics for objects.

Table 8-160. VMware Cloud on AWS Metrics

Object Type	Metric Key	Metric Value	Description
Bill	Cost Monthly Commit Expense	Double	Represents the total amount spent on the Commit purchases for a month.
	Cost Monthly OnDemand Expense	Double	Represents the total amount spent on the OnDemand purchases for a month.
	Cost Monthly Total Expense	Double	Represents the total amount spent on the OnDemand and Commit purchases for a month.

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	Cost Outstanding Expense	Double	Represents the daily Outstanding expenses.
Component	Cost Component Expense	Double	Represents the amount spent for the purchases of Commit or OnDemand components for a month.
Org Object	Configuration Maximum Number of hosts per Organization Soft Limit	Double	Represents the number of hosts per organization.
	Configuration Maximum Number of hosts per Organization Provisioned	Double	
	Configuration Maximum Number of hosts per Organization Soft Limit % Used	Double	
	Configuration Maximum Public IP Addresses (Elastic IPs) Soft Limit	Double	Represents the maximum number of IP addresses per organization.
	Configuration Maximum Public IP Addresses (Elastic IPs) Provisioned	Double	
	Configuration Maximum Public IP Addresses (Elastic IPs) Soft Limit % Used	Double	
	Configuration Maximum Number of SDDCs per Organization Soft Limit	Double	Represents the maximum number of SDDCs per organization.
	Configuration Maximum Number of SDDCs per Organization Provisioned Limit	Double	
	Configuration Maximum Number of SDDCs per Organization Soft Limit % Used	Double	
SDDC	VMC Configuration Maximums Linked VPC Count Limit	Double	Represents the maximum number of linked AWS VPCs per SDDC.
	VMC Configuration Maximums Linked VPC Count Provisioned	Double	
	VMC Configuration Maximums Linked VPC Count Limit % Used	Double	
	Configuration Maximum Max clusters Soft Limit	Double	Represents the maximum number of vSphere clusters per SDDC.
	Configuration Maximum Max clusters Hard Limit	Double	
	Configuration Maximum Max clusters Provisioned	Double	
	Configuration Maximum Max clusters Soft Limit % Used	Double	

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	Configuration Maximum Max clusters Hard Limit % Used	Double	Represents the maximum number of ESXi hosts per SDDC.
	Configuration Maximum Maximum hosts per SDDC Limit	Double	
	Configuration Maximum Maximum hosts per SDDC Provisioned	Double	
	Configuration Maximum Maximum hosts per SDDC Limit % Used	Double	
	Configuration Maximum Maximum VMs per SDDC Limit	Double	Represents the maximum number of virtual machines per SDDC.
	Configuration Maximum Maximum VMs per SDDC Provisioned	Double	
	Configuration Maximum Maximum VMs per SDDC Limit % Used	Double	
	VMC Configuration Maximums MGW Gateway Firewall Rule Count Limit	Double	Represents the maximum number of Management Gateway Firewall rules.
	VMC Configuration Maximums MGW Gateway Firewall Rule Count Provisioned	Double	
	VMC Configuration Maximums MGW Gateway Firewall Rule Count Limit % Used	Double	
	VMC Configuration Maximums CGW Gateway Firewall Rule Count Limit	Double	Represents the maximum number of Compute Gateway Firewall rules.
	VMC Configuration Maximums CGW Gateway Firewall Rule Count Provisioned	Double	
	VMC Configuration Maximums CGW Gateway Firewall Rule Count Limit % Used	Double	
	VMC Configuration Maximums Direct Connect private VIF Connection Count Limit	Double	Represents the maximum number of private virtual interfaces attached to one SDDC.
	VMC Configuration Maximums Direct Connect private VIF Connection Count Provisioned	Double	
	VMC Configuration Maximums Direct Connect private VIF Connection Count Limit % Used	Double	
Cluster Compute Resource	Configuration Maximum Min hosts per cluster for full SLA Status	Double	Represents the minimum number of ESXi per vSphere cluster that must be supported at full SLA.

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	Configuration Maximum Minimum hosts per cluster for full SLA Limit Violated	Double	Represents the minimum number of ESXi hosts per vSphere cluster with no SLA.
	Configuration Maximum Min hosts per cluster for no SLA Limit	Double	
	Configuration Maximum Min hosts per cluster for no SLA Limit Violated	Double	
	Configuration Maximum Max hosts per cluster (including stretched clusters) Limit	Double	Represents the maximum number of ESXi hosts per vSphere cluster. This limit applies to both single-AZ clusters and stretched clusters.
	Configuration Maximum Max hosts per cluster (including stretched clusters) Provisioned	Double	
	Configuration Maximum Max hosts per cluster (including stretched clusters) Limit % Used	Double	
Resource Pool	CPU vCPUs Allocated to all Consumers	Double	Represents the number of vCPUs allocated to the vCenter and NSX management appliances in a regular-sized SDDC.
	Memory Memory Allocated to all Consumers	Double	Represents the RAM allocated to the vCenter and NSX management appliances in a large and regular sized SDDC.
Host System	Configuration Maximum VMs per host Limit	Double	Represents the maximum number of VMs per host.
	Summary Total Number of VMs	Double	
	VMC Configuration Maximum VMs per host Limit % Used	Double	
Logical Router	VMC Configuration Maximums IPSec VPN Tunnel Count Limit	Double	Represents the maximum number of IPSec VPN tunnels created per SDDC.
	VMC Configuration Maximums IPSec VPN Tunnel Count Provisioned	Double	
	VMC Configuration Maximums IPSec VPN Tunnel Count Limit % Used	Double	
	VMC Configuration Maximums L2VPN Client Count Limit	Double	Represents the maximum number of sites connecting to L2 VPN server per SDDC.
	VMC Configuration Maximums L2VPN Client Count Provisioned	Double	
	VMC Configuration Maximums L2VPN Client Count Limit % Used	Double	

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
Logical Switch	VMC Configuration Maximums Logical Segment Count Limit	Double	Represents the maximum number of logical segments per SDDC.
	VMC Configuration Maximums Logical Segment Count Provisioned	Double	
	VMC Configuration Maximums Logical Segment Count Limit % Used	Double	
	VMC Configuration Maximums Logical Ports Count Limit	Double	Represents the maximum number of ports on a logical segment.
	VMC Configuration Maximums Logical Ports Count Provisioned	Double	
	VMC Configuration Maximums Logical Ports Count Limit % Used	Double	
	VMC Configuration Maximums Extended Network Count Limit	Double	Represents the maximum number of logical segments extended from on-premises.
	VMC Configuration Maximums Extended Network Count Provisioned	Double	
	VMC Configuration Maximums Extended Network Count Limit % Used	Double	
Router Service (NAT Rules)	VMC Configuration Maximums NAT Rule Count Limit	Double	Represents the maximum number of Compute Gateway NAT rules.
	VMC Configuration Maximums NAT Rule Count Provisioned	Double	
	VMC Configuration Maximums NAT Rule Count Limit % Used	Double	
Group	VMC Configuration Maximums Distributed Firewall Grouping Object Count Limit	Double	Represents the maximum number of grouping objects (security groups).
	VMC Configuration Maximums Distributed Firewall Grouping Object Count Provisioned	Double	
	VMC Configuration Maximums Distributed Firewall Grouping Object Count Limit % Used	Double	
	VMC Configuration Maximums IP Address Count Limit	Double	Represents the maximum number of IP addresses that can be included in an IP set.
	VMC Configuration Maximums IP Address Count Provisioned	Double	
	VMC Configuration Maximums IP Address Count Limit % Used	Double	

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	VMC Configuration Maximums Distributed Firewall Rule Count Limit	Double	Represents the maximum number of distributed firewall rules per grouping object (security group).
	VMC Configuration Maximums Distributed Firewall Rule Count Provisioned	Double	
	VMC Configuration Maximums Distributed Firewall Rule Count Limit % Used	Double	
	VMC Configuration Maximums vm Count Limit	Double	Represents the maximum number of VMs per grouping object (security group).
	VMC Configuration Maximums vm Count Provisioned	Double	
	VMC Configuration Maximums vm Count Limit % Used	Double	
Firewall Sections	VMC Configuration Maximums Distributed Firewall Section Count Limit	Double	Represents the maximum number of distributed firewall sections.
	VMC Configuration Maximums Distributed Firewall Section Count Provisioned	Double	
	VMC Configuration Maximums Distributed Firewall Section Count Limit % Used	Double	
	VMC Configuration Maximums Distributed Firewall Rule Count Limit	Double	Represents the maximum number of distributed firewall rules across all sections groups such as, Emergency Rules, Infrastructure Rules, and so on.
	VMC Configuration Maximums Distributed Firewall Rule Count Provisioned	Double	
	VMC Configuration Maximums Distributed Firewall Rule Count Limit % Used	Double	
	VMC Configuration Maximums Distributed (Group_Name) Firewall Rule Count Limit	Double	Represents the maximum number of distributed firewall rules per section group.
	VMC Configuration Maximums Distributed (Group_Name) Firewall Rule Count Provisioned	Double	
	VMC Configuration Maximums Distributed (Group_Name) Firewall Rule Count Limit % Used	Double	
	VMC Configuration Maximums Distributed (Group_Name) Firewall Section Count Limit	Double	Represents the maximum number of distributed firewall sections per section group, such as,

Table 8-160. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	VMC Configuration Maximums Distributed (Group_Name) Firewall Section Count Provisioned	Double	Emergency Rules, Infrastructure Rules, and so on.
	VMC Configuration Maximums Distributed (Group_Name) Firewall Section Count Limit % Used	Double	
Virtual Machine	VMC Configuration Maximums Security Tag Count Limit	Double	Represents the maximum number of security tags per VM.
	VMC Configuration Maximums Security Tag Count Provisioned	Double	
	VMC Configuration Maximums Security Tag Count Limit % Used	Double	
Management Cluster	VMC Configuration Maximums IPFIX Collector Count Limit	Double	Represents the maximum number of IPFIX Collectors configured.
	VMC Configuration Maximums IPFIX Collector Count Provisioned	Double	
	VMC Configuration Maximums IPFIX Collector Count Limit % Used	Double	
Datastore	Configuration Maximum Maximum datastore capacity that can be utilized Limit	Double	Represents the maximum datastore capacity that can be utilized. You can use up to 75% of available datastore capacity. Usage beyond this point creates a non-compliant environment as described in Service Level Agreement for VMware Cloud on AWS .
	Configuration Maximum Datastore capacity requiring remediation plan Limit	Double	Represents the datastore capacity that requires a remediation plan. You must prepare a remediation plan when capacity utilization nears 70%. You can either add hosts to augment datastore capacity or reduce storage utilization.

Table 8-161. VMware Cloud on AWS Metrics Properties

Object Type	Property Name	Property Value	Description
Bill	Configuration Currency	String	Represents the currency unit set in the VMware Cloud on AWS account by the customer.
	Configuration OrgId	String	Represents the organization ID for the associated bill.

Table 8-161. VMware Cloud on AWS Metrics Properties (continued)

Object Type	Property Name	Property Value	Description
	Configuration Statement Bill Start Date	String	Represents the start date of the statement bill.
	Configuration Statement Bill End Date	String	Represents the end date of the statement bill.
	Summary YTD Commit Expense	Double	Represents the total amount spent on the Commit purchases for the current calendar year until the last generated statement bill.
	Summary YTD OnDemand Expense	Double	Represents the total amount spent on the OnDemand purchases for the current calendar year until the last generated statement bill.
	Summary YTD Total Expense	Double	Represents the total amount spent on the Commit and OnDemand purchases for the current calendar year until the last generated statement bill.
Component	Configuration Component Start Date	String	Represents the billing start date of the component purchase.
	Configuration Component End Date	String	Represents the billing end date of the component purchase.
	Configuration Component SKU Description	String	Represents the SKU of the component.
	Configuration Component Service Type	String	Represents the component service type.
	Configuration Component Usage Type	String	Represents the component usage type.
	Configuration Subscription Status	boolean	Represents whether a Commit is still available for use.
	Summary Number of Units Used	Integer	Represents the total number of components.
Org	Configuration Id	String	Represents the organization ID.
	Configuration Name	String	Represents the organization name.

Metrics in NSX-T Adapter

The NSX-T adapter collects metrics for objects within its plug-in.

Table 8-162. Metrics in the NSX-T On-Premise

Resource	Metrics	Metric Keys
Management Cluster	System Capacity <ul style="list-style-type: none"> ■ Max Supported Count ■ Max Threshold Percentage ■ Min Threshold Percentage ■ Usage Count ■ Usage Count Percentage ■ Severity 	System Capacity Keys <ul style="list-style-type: none"> ■ System Capacity <Object_Kind> MaxSupportedCount ■ System Capacity <Object_Kind> MaxThresholdPercentage ■ System Capacity <Object_Kind> MinThresholdPercentage ■ System Capacity <Object_Kind> UsageCount ■ System Capacity <Object_Kind> UsageCountPercentage ■ System Capacity <Object_Kind> Severity
Transport Node	<ul style="list-style-type: none"> ■ CPU <ul style="list-style-type: none"> ■ CPU Cores ■ DPDK CPU Cores ■ DPDK CPU Core Average Usage ■ DPDK CPU Core Highest Usage ■ Non-DPDK CPU Core Average Usage ■ Non-DPDK CPU Core Highest Usage ■ Memory <ul style="list-style-type: none"> ■ Total ■ Used ■ Cache ■ Total Swap ■ Used Swap 	<ul style="list-style-type: none"> ■ CPU Metric Keys <ul style="list-style-type: none"> ■ Cpu Cores ■ Cpu DPDKCores ■ Cpu AvgDpdkCpuCoreUsage ■ Cpu HighDpdkCpuCoreUsage ■ Cpu AvgNonDpdkCpuCoreUsage ■ Cpu HighNonDpdkCpuCoreUsage ■ Memory metric keys <ul style="list-style-type: none"> ■ Memory Total ■ Memory Used ■ Memory Cache ■ Memory Total Swap ■ Memory Used Swap
	File Systems <FileSystemMount> Used	FileSystems Used
	Statistics Interface <InterfaceID> <ul style="list-style-type: none"> ■ Received Data (bytes) ■ Received Packets dropped ■ Received Packets errors ■ Received Framing errors ■ Received Packets ■ Transmitted Data (bytes) ■ Transmitted Packets dropped ■ Transmitted Packets errors ■ Transmitted carrier losses detected ■ Transmitted Packets ■ Transmitted Collisions detected 	Statistics Metric Keys <ul style="list-style-type: none"> ■ stats Interface RxDData ■ stats Interface RxDropped ■ stats Interface RxEErrors ■ stats Interface RxFFrame ■ stats Interface RXPackets ■ stats Interface TxData ■ stats Interface TxDropped ■ stats Interface TxErrors ■ stats Interface TxCarrier ■ stats Interface TxPackets ■ stats Interface TxColls

Table 8-162. Metrics in the NSX-T On-Premise (continued)

Resource	Metrics	Metric Keys
Load Balancer Service	<ul style="list-style-type: none"> ■ CPU Usage(%) ■ Memory Usage(%) ■ Active Transport Nodes ■ Standby Transport Nodes ■ Sessions: <ul style="list-style-type: none"> ■ L4Average ■ L4Current ■ L4Maximum ■ L4Total ■ L7Average ■ L7Current ■ L7Maximum ■ L7Total 	<ul style="list-style-type: none"> ■ CPU Usage ■ Memory Usage ■ Active Transport Nodes ■ Standby Transport Nodes ■ Sessions L4Average ■ Sessions L4Current ■ Sessions L4Maximum ■ Sessions L4Total ■ Sessions L7Average ■ Sessions L7Current ■ Sessions L7Maximum ■ Sessions L7Total
Load Balancer Virtual Server	<ul style="list-style-type: none"> ■ Statistics <ul style="list-style-type: none"> ■ Bytes Inbound Bytes Total ■ Bytes Average Inbound Bytes Per Second ■ Bytes Outbound Bytes Total ■ Bytes Average Outbound Bytes Per Second ■ Http Http Request Rate ■ Http Http Requests ■ Packets Inbound Packets Total ■ Packets Inbound Packets Rate ■ Packets Outbound Packets Total ■ Packets Outbound Packets Rate ■ Packets Dropped ■ Sessions <ul style="list-style-type: none"> ■ Average Current Sessions Per Second ■ Current Sessions ■ Maximum Sessions ■ Dropped Sessions ■ Total Sessions 	<ul style="list-style-type: none"> ■ Statistics metric keys <ul style="list-style-type: none"> ■ stats Bytes Inbound ■ stats Bytes InboundRate ■ stats Bytes Outbound ■ stats Bytes OutboundRate ■ stats Http RequestRate ■ stats Http Requests ■ stats Packets Inbound ■ stats Packets InboundRate ■ stats Packets Outbound ■ stats Packets OutboundRate ■ stats Packets Dropped ■ Sessions metric keys <ul style="list-style-type: none"> ■ Sessions CurrentRate ■ Sessions Current ■ Sessions Maximum ■ Sessions Dropped ■ Sessions Total

Table 8-162. Metrics in the NSX-T On-Premise (continued)

Resource	Metrics	Metric Keys
Load Balancer Pool	<ul style="list-style-type: none"> ■ Statistics <ul style="list-style-type: none"> ■ Bytes Inbound Bytes Total ■ Bytes Average Inbound Bytes Per Second ■ Bytes Outbound Bytes Total ■ Bytes Average Outbound Bytes Per Second ■ Http Http Request Rate ■ Http Http Requests ■ Packets Inbound Packets Total ■ Packets Inbound Packets Rate ■ Packets Outbound Packets Total ■ Packets Outbound Packets Rate ■ Packets Dropped ■ Sessions <ul style="list-style-type: none"> ■ Average Current Sessions Per Second ■ Current Sessions ■ Maximum Sessions ■ Dropped Sessions ■ Total Sessions 	<ul style="list-style-type: none"> ■ Statistics metric keys <ul style="list-style-type: none"> ■ stats Bytes Inbound ■ stats Bytes InboundRate ■ stats Bytes Outbound ■ stats Bytes OutboundRate ■ stats HttpRequestRate ■ stats HttpRequests ■ stats Packets Inbound ■ stats Packets InboundRate ■ stats Packets Outbound ■ stats Packets OutboundRate ■ stats Packets Dropped ■ Sessions Metric metric keys <ul style="list-style-type: none"> ■ Sessions CurrentRate ■ Sessions Current ■ Sessions Maximum ■ Sessions Dropped ■ Sessions Total
Management Services	<ul style="list-style-type: none"> ■ Service Monitor Process ID ■ Service Monitor Runtime state ■ Service Process ID ■ Service Runtime State 	<ul style="list-style-type: none"> ■ ServiceMonitorProcessId ■ ServiceMonitorRuntimeState ■ ServiceProcessIds ■ ServiceRuntimeState
Logical Router	Statistics <ul style="list-style-type: none"> ■ Received Data (bytes) ■ Received Packets dropped ■ Received Packets ■ Transmitted Data (bytes) ■ Transmitted Packets dropped ■ Transmitted Packets 	Statistics metric keys <ul style="list-style-type: none"> ■ stats RxDData ■ stats RxDropped ■ stats RxDPackets ■ stats TxData ■ stats TxDropped ■ stats TxPackets

Table 8-162. Metrics in the NSX-T On-Premise (continued)

Resource	Metrics	Metric Keys
	Configuration Maximums <ul style="list-style-type: none"> Router Port Count ARP Entries Count Tier 1 Router Count Route Map Count Route Maps <RouteMapName:RouteMapId> Rule Count Prefix List Count IP Prefix Lists <IPPrefixListName:IPPrefixListId> Prefix List Entries Count 	Configuration Maximums metric keys <ul style="list-style-type: none"> configMax routerPortCount configMax routerArpEntryCount <hr/> Note Metric applicable for T1 router. <ul style="list-style-type: none"> configMax tier1RouterCount configMax routeMapCount configMax RouteMaps routeMapRuleCount <hr/> Note Metric applicable for T0 router. <ul style="list-style-type: none"> configMax prefixListCount configMax IPPrefixLists prefixListEntriesCount <hr/> Note Metric applicable for T0 and T1 router.
Logical Switch	Statistics <ul style="list-style-type: none"> Inbound Bytes Total Inbound Bytes Dropped Inbound Bytes Throughput Outbound Bytes Total Outbound Bytes Dropped Outbound Bytes Throughput Inbound Packets Total Inbound Packets Dropped Inbound Packets Throughput Outbound Packets Total Outbound Packets Dropped Outbound Packets Throughput 	Metric keys <ul style="list-style-type: none"> stats IngressBytes stats IngressBytesDropped stats IngressBytesThroughput stats IngressPackets stats IngressPacketsDropped stats IngressPacketsThroughput stats EgressBytes stats EgressBytesDropped stats EgressBytesThroughput stats EgressPackets stats EgressPacketsDropped stats EgressPacketsThroughput
Logical Switch Group	Configuration Maximums <ul style="list-style-type: none"> Logical Segment Count 	Metric keys <ul style="list-style-type: none"> configMax LogicalSegmentCount
Management Appliances	Management Node Count	Management node count
Manager Node	<ul style="list-style-type: none"> File Systems <FileSystemMount> <ul style="list-style-type: none"> File System Id File System Type Total (KB) Used(KB) Used(%) 	File Systems Metric Keys <ul style="list-style-type: none"> FileSystems <FileSystemMount> FileSystemId FileSystems <FileSystemMount> Type FileSystems <FileSystemMount> Total FileSystems <FileSystemMount> Used FileSystems <FileSystemMount> usedPercentage

Table 8-162. Metrics in the NSX-T On-Premise (continued)

Resource	Metrics	Metric Keys
	Network Interfaces <InterfaceID> <ul style="list-style-type: none"> Received Data Bits per second Received Data Cumulative(bytes) Received Framing Errors Cumulative Received Framing Errors Per second Received Packets Cumulative Received Packets Per Second Received Packets Dropped Cumulative Received Packets Dropped Per second Received Packets Error Cumulative Received Packets Error Per second Transmitted Carrier losses detected Cumulative Transmitted Carrier losses detected Per second Transmitted Collisions detected Cumulative Transmitted Collisions detected Per second Transmitted Data Bits per second Transmitted Data Cumulative(bytes) Transmitted Packets Cumulative Transmitted Packets Per second Transmitted Packets Dropped Cumulative Transmitted Packets Dropped Per second Transmitted Packets errors Cumulative Transmitted Packets errors Per second 	Network Interface metric keys <ul style="list-style-type: none"> Interfaces <InterfaceID> RxData BitsPerSecond Interfaces <InterfaceID> RxData Cumulative Interfaces <InterfaceID> RxFrame Cumulative Interfaces <InterfaceID> RxFrame PerSecond Interfaces <InterfaceID> RxPackets Cumulative Interfaces <InterfaceID> RxPackets PerSecond Interfaces <InterfaceID> RxDropped Cumulative Interfaces <InterfaceID> RxDropped PerSecond Interfaces <InterfaceID> RxErrors Cumulative Interfaces <InterfaceID> RxErrors PerSecond Interfaces <InterfaceID> TxCarrier Cumulative Interfaces <InterfaceID> TxCarrier PerSecond Interfaces <InterfaceID> TxColls Cumulative Interfaces <InterfaceID> TxColls PerSecond Interfaces <InterfaceID> TxData BitsPerSecond Interfaces <InterfaceID> TxData Cumulative Interfaces <InterfaceID> TxPackets Cumulative Interfaces <InterfaceID> TxPackets PerSecond Interfaces <InterfaceID> TxDropped Cumulative Interfaces <InterfaceID> TxDropped PerSecond Interfaces <InterfaceID> TxErrors Cumulative Interfaces <InterfaceID> TxErrors PerSecond
	CPU <ul style="list-style-type: none"> CPU Cores DPDK CPU Cores DPDK CPU Core Average Usage DPDK CPU Core Highest Usage Non-DPDK CPU Core Average Usage Non-DPDK CPU Core Highest Usage 	CPU Metric Keys <ul style="list-style-type: none"> Cpu Cores Cpu DPDKCores Cpu AvgDpdkCpuCoreUsage Cpu HighDpdkCpuCoreUsage Cpu AvgNonDpdkCpuCoreUsage Cpu HighNonDpdkCpuCoreUsage
	Memory <ul style="list-style-type: none"> Total Used Cache Total Swap Used Swap 	Memory metric keys <ul style="list-style-type: none"> Memory Total Memory Used Memory Cache Memory TotalSwap Memory UsedSwap

Table 8-162. Metrics in the NSX-T On-Premise (continued)

Resource	Metrics	Metric Keys
Controller Cluster	<ul style="list-style-type: none"> ■ Controller Node Count ■ Cluster Status Controller Cluster Status ■ Cluster Status Management cluster Status 	<p>Controller cluster metrics keys</p> <ul style="list-style-type: none"> ■ Cluster Status Controller Node Count ■ ClusterStatus ControllerClusterStatus ■ ClusterStatus ManagementClusterStatus <p>Note These metrics are not collected for NSX-T version above 2.4</p>
Controller Node	<ul style="list-style-type: none"> ■ Connectivity Status Cluster Connectivity ■ Connectivity Status Manager Connectivity ■ File System ID ■ File System Type ■ Total(KB) ■ Used(KB) ■ Used(%) ■ Network Interfaces <InterfaceID> ■ Received Data Bits per second ■ Received Data Cumulative(bytes) ■ Received Framing Errors Cumulative ■ Received Framing Errors Per second ■ Received Packets Cumulative ■ Received Packets Per Second ■ Received Packets Dropped Cumulative ■ Received Packets Dropped Per second ■ Received Packets Error Cumulative ■ Received Packets Error Per second ■ Transmitted Carrier losses detected Cumulative ■ Transmitted Carrier losses detected Per second ■ Transmitted Collisions detected Cumulative ■ Transmitted Collisions detected Per second ■ Transmitted Data Bits per second ■ Transmitted Data Cumulative(bytes) ■ Transmitted Packets Cumulative ■ Transmitted Packets Per second ■ Transmitted Packets Dropped Cumulative ■ Transmitted Packets Dropped Per second ■ Transmitted Packets errors Cumulative ■ Transmitted Packets errors Per second 	<p>Note These metrics are not collected for NSX-T version above 2.4</p> <ul style="list-style-type: none"> ■ ConnectivityStatus ClusterConnectivity ■ ConnectivityStatus ManagerConnectivity ■ FileSystems <FileSystemMount> FileSystemId ■ FileSystems <FileSystemMount> Type ■ FileSystems <FileSystemMount> Total ■ FileSystems <FileSystemMount> Used ■ FileSystems <FileSystemMount> usedPercentage ■ Interfaces <InterfaceID> RxData BitsPerSecond ■ Interfaces <InterfaceID> RxData Cumulative ■ Interfaces <InterfaceID> RxFrame Cumulative ■ Interfaces <InterfaceID> RxFrame PerSecond ■ Interfaces <InterfaceID> RxPackets Cumulative ■ Interfaces <InterfaceID> RxPackets PerSecond ■ Interfaces <InterfaceID> RxDropped Cumulative ■ Interfaces <InterfaceID> RxDropped PerSecond ■ Interfaces <InterfaceID> RxErrors Cumulative ■ Interfaces <InterfaceID> RxErrors PerSecond ■ Interfaces <InterfaceID> TxCarrier Cumulative ■ Interfaces <InterfaceID> TxCarrier PerSecond ■ Interfaces <InterfaceID> TxColls Cumulative ■ Interfaces <InterfaceID> TxColls PerSecond ■ Interfaces <InterfaceID> TxData BitsPerSecond ■ Interfaces <InterfaceID> TxData Cumulative ■ Interfaces <InterfaceID> TxPackets Cumulative ■ Interfaces <InterfaceID> TxPackets PerSecond ■ Interfaces <InterfaceID> TxDropped Cumulative ■ Interfaces <InterfaceID> TxDropped PerSecond ■ Interfaces <InterfaceID> TxErrors Cumulative ■ Interfaces <InterfaceID> TxErrors PerSecond

Table 8-163. Metrics in the NSX-T on VMware Cloud on AWS

Resource	Metrics	Metric Keys
Logical Router	<p>The following metrics are specify to Tier 0 Router.</p> <p>Statistics Interface</p> <ul style="list-style-type: none"> ■ Received Data (Bytes) ■ Received Packets ■ Received Packets Dropped ■ Transmitted Data ■ Transmitted Received Data (Bytes) ■ Transmitted Received Packets ■ Transmitted Received Packets Dropped 	<p>Stats Metrics</p> <p>Statistics Interface</p> <ul style="list-style-type: none"> ■ stats Interface RxDData ■ stats Interface RxDPackets ■ stats Interface RxDropped ■ stats Interface TxData ■ stats Interface TxPackets ■ stats Interface TxDropped <hr/> <p>Note These metrics are only for Tier 0 Router.</p>
Firewall Section Group	<p>Configuration Maximums</p> <ul style="list-style-type: none"> ■ Distributed Firewall Section Count ■ Distributed Firewall Rule Count ■ MGW Gateway Firewall Rule Count ■ CGW Gateway Firewall Rule Count ■ Distributed Application Firewall Rule Count ■ Distributed Application Firewall Section Count ■ Distributed Environment Firewall Rule Count ■ Distributed Environment Firewall Section Count ■ Distributed Infrastructure Firewall Rule Count ■ Distributed Infrastructure Firewall Section Count ■ Distributed Emergency Firewall Rule Count ■ Distributed Emergency Firewall Section Count ■ Distributed Ethernet Firewall Rule Count 	<p>Configuration metric keys</p> <ul style="list-style-type: none"> ■ configMax MaxDistributedFirewallSections ■ configMax MaxDistributedFirewallRules ■ configMax MaxMGWGatewayFirewallRules ■ configMax MaxCGWGatewayFirewallRules ■ configMax MaxDistributedApplicationFirewallRules ■ configMax MaxDistributedApplicationFirewallSections ■ configMax MaxDistributedEnvironmentFirewallRules ■ configMax MaxDistributedEnvironmentFirewallSections ■ configMax MaxDistributedInfrastructureFirewallRules ■ configMax MaxDistributedInfrastructureFirewallSections ■ configMax MaxDistributedEmergencyFirewallRules ■ configMax MaxDistributedEmergencyFirewallSections ■ configMax MaxDistributedEthernetFirewallRules ■ configMax MaxDistributedEthernetFirewallSections <hr/> <p>Note These metrics are only for NSX-T on VMware Cloud on AWS. For NSX-T on-premise, the values for these metrics is shown as zero.</p>

Table 8-163. Metrics in the NSX-T on VMware Cloud on AWS (continued)

Resource	Metrics	Metric Keys
	<ul style="list-style-type: none"> ■ Distributed Ethernet Firewall Section Count <hr/> <p>Note These metrics are only for NSX-T on VMware Cloud on AWS. For NSX-T on-premise, the values for these metrics show zero.</p>	
Logical Switch Group	Configuration Maximums <ul style="list-style-type: none"> ■ Logical Segment Count ■ Extended Network Count 	Metric Keys <ul style="list-style-type: none"> ■ configMax LogicalSegmentCount ■ configMax ExtendedNetworkcount <hr/> <p>Note The metric (configMax ExtendedNetworkcount) is only for NSX-T on VMware Cloud on AWS. For NSX-T on-premise, its value is zero.</p>

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 Alert 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.

- [vSphere Pod Alert Definitions](#)

The vCenter adapter provides alert definitions that generate alerts on the vSphere Pod objects in your environment.

- [VMware Cloud on AWS 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 **VMware Cloud on AWS** objects.

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.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 1 Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level. 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 right sizing of VMs.
Fully-automated DRS-enabled cluster has CPU contention caused by more than half of the virtual machines.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 1 Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level. 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 right sizing 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: <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 1 Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level. 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 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 right sizing of VMs.
Fully-automated DRS-enabled cluster has high CPU workload.	Symptoms include all of the following: <ul style="list-style-type: none"> ■ DRS enabled ■ DRS fully automated ■ Cluster CPU workload above DT ■ Cluster CPU workload at warning/immediate/critical level 	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 1 Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level. 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 To increase memory capacity add more hosts to the cluster. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for right sizing of VMs.

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has memory contention caused by more than half of the virtual machines.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 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 right sizing of VMs.
Fully-automated DRS-enabled cluster has memory contention caused by overpopulation of virtual machines.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 1 Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level. 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 right sizing 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.	<ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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 high memory workload and contention.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 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.
vSphere High Availability (HA) failover resources are insufficient	vSphere High Availability (HA) failover resources are insufficient	<p>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.</p>

Alert Definition	Symptoms	Recommendations
vSphere HA master missing.	vCenter Server is unable to find a master vSphere HA agent (fault symptom)	
Proactive HA provider has reported health degradation on the underlying hosts.	Proactive HA provider reported host health degradation.	Contact your hardware vendor support.

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

Standalone host has CPU contention caused by overpopulation of virtual machines.

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: <ul style="list-style-type: none"> ■ 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 <ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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 overpopulation of virtual machines.	Symptoms include the following: <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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 caused by overpopulation of virtual machines.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ [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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ Host network received packets dropped ■ Host network transmitted packets dropped 	<ol style="list-style-type: none"> 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.
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.	<p>Symptoms include the following:</p> <ul style="list-style-type: none"> ■ 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.

Alert Definition	Symptoms	Recommendations
Baseboard Management Controller sensors are reporting problems.	Symptoms include the following: <ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.
Path redundancy to storage device degraded	<ul style="list-style-type: none"> ■ A path to storage device went down ■ Host has no redundancy to storage device 	See KB topic, <i>Path redundancy to the storage device is degraded</i> (1009555)
Power sensors are reporting problems.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.

Alert Definition	Symptoms	Recommendations
SEL sensors are reporting problems.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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 primary agent is able to connect to the HA agent on the host. Note You can use the Duplicate IP warning in the <code>/var/log/vmkernel</code> log file on an ESX host or the <code>/var/log/messages</code> log file on an ESXi host to identify the computer that has the duplicate IP address.
Host is experiencing network contention caused by too much traffic.	Symptoms include all the following: <ul style="list-style-type: none"> ■ Host is experiencing dropped network packets ■ Host network workload at warning/immediate/critical level 	<ol style="list-style-type: none"> 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).	<p>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:</p> <ul style="list-style-type: none"> ■ 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: <code>esxconfig-nics</code>. The output that appears is similar to the following: <pre>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 <code>esxconfig-nics -l</code>", and observe the affected vmnics.</pre> <p>Verify that the vmnic identified in the alert is still connected to the switch and configured properly:</p> <ul style="list-style-type: none"> ■ 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.

Alert Definition	Symptoms	Recommendations
		<ul style="list-style-type: none"> ■ Check for activity between the physical switch and the vmnic. You can check activity by performing a network trace or observing activity LEDs. ■ Check for network port settings on the physical switch. <p>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.</p>
The host lost connectivity to a Network File System (NFS) server.	Lost connection to NFS server (fault symptom).	<ol style="list-style-type: none"> 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 <code>vmkping</code> to ping the NFS server: <code>"vmkping <nfs server>"</code>. 5 For advanced troubleshooting information, see http://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).	<p>To determine the actual failure or to eliminate possible problems, first connect to ESX through SSH or the console:</p> <ol style="list-style-type: none"> 1 Identify the available uplinks by running <code>esxcfg-nics -l</code>. 2 Remove the reported vmnic from the port groups by running <code>esxcfg-vswitch -U <affected vmnic> affected vSwitch</code>. 3 Link available uplinks to the affected port groups by running <code>esxcfg-vswitch -L <available vmnic> affected vSwitch</code>. <p>Next, check the status of the vmnic in vSphere Client or the ESX service console:</p> <ol style="list-style-type: none"> 1 In vSphere Client, select the ESX host, click Configuration, and then click Networking. <p>The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently unavailable.</p> <ol style="list-style-type: none"> 2 From the service console, run <code>esxcfg-nics -l</code>. The output that appears is similar to the following example: Name PCI Driver Link Speed Duplex Description. <pre> ----- ----- 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 </pre>

Alert Definition	Symptom	Recommendations
		<p>to the intended physical switch ports. To verify the connections, shut down each 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:</p> <ol style="list-style-type: none"> 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. <p>If the problem is caused by hardware, contact your hardware vendor for a hardware replacement.</p>
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.	<ul style="list-style-type: none"> ■ 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 Alert 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 for the vSAN Cluster Object

Alerts on the vSAN Cluster object have health, risk, and efficiency impact.

Table 8-164. vSAN Cluster Object Health Alert Definitions

Alert	Alert Type	Alert Subtype	Description
Basic (unicast) connectivity check (normal ping) has failed on vSAN host.	Storage	Configuration	Triggered when basic (unicast) connectivity check (normal ping) has failed on the vSAN host due to network misconfiguration.
Check the free space on physical disks in the vSAN cluster.	Storage	Availability	Triggered when a check of free space on physical disks in the vSAN cluster results in an error or warning.
CLOMD process on the host has issues and impacting the functionality of vSAN cluster.	Storage	Availability	Triggered when CLOMD process on the host has issues and impacting the functionality of vSAN cluster.

Table 8-164. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
Disk load variance between some vSAN disks exceeded the threshold value.	Storage	Performance	Triggered when disk load variance between some vSAN disks exceeded the threshold value. vSAN cannot perform the load balance properly.
Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster.	Storage	Configuration	Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster.
Host has invalid unicast agent and impacting the health of vSAN Stretched Cluster.	Storage	Configuration	Triggered when the 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.
Host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic.	Network	Configuration	Triggered when the 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.
Host in a vSAN cluster has connectivity issues and vCenter Server does not know its state.	Network	Configuration	Triggered when the host in a vSAN cluster has connectivity issues and vCenter Server does not know its state.
Host in a vSAN cluster has IP multicast connectivity issue.	Network	Configuration	Triggered when the 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.
Host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host.	Storage	Configuration	Triggered when the host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host.
Network latency check of vSAN hosts failed. It requires < 1 ms RTT.	Network	Configuration	Triggered if network latency check of vSAN hosts is greater than or equal to 1 ms RTT.
One or more hosts in the vSAN cluster have misconfigured multicast addresses.	Network	Configuration	Triggered when one or more hosts in the vSAN cluster have misconfigured multicast addresses.
One or more physical disks on vSAN host is experiencing software state health issues.	Storage	Availability	Triggered when one or more physical disks on vSAN host is experiencing software state health issues.
One or more vSAN enabled hosts are not in the same IP subnet.	Network	Configuration	Triggered when one or more vSAN enabled hosts are not in the same IP subnet.

Table 8-164. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
Overall health of the physical disks in a vSAN Cluster is impacted.	Storage	Availability	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.
Overall health of VMs residing on vSAN datastore is reporting issues.	Storage	Availability	Triggered when overall health of the VMs on a vSAN datastore is impacted.
Overall health of vSAN objects is reporting issues.	Storage	Availability	Triggered when overall health of vSAN objects is reporting issues.
Ping test with large packet size between all VMKernel adapters with vMotion traffic enabled has issues.	Network	Configuration	Triggered when ping test with large packet size between all VMKernel adapter with vMotion traffic enabled is impacted.
Ping test with small packet size between all VMkernel adapters with vMotion traffic enabled has issues.	Network	Configuration	Triggered when ping test with small packet size between all VMKernel adapter with vMotion traffic enabled is impacted.
Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster.	Storage	Performance	Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster.
Statistics collection of vSAN performance service is not working correctly.	Storage	Availability	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.
MTU check (ping with large packet size) has failed on vSAN host.	Storage	Configuration	Triggered when MTU check (ping with large packet size) has failed on vSAN environment due to some MTU misconfiguration in the vSAN network.
The preferred fault domain is not set for the witness host in a vSAN Stretched cluster.	Storage	Configuration	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.
Unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster.	Storage	Configuration	Triggered when unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster.
vCenter Server has lost connection to a host that is part of a vSAN cluster.	Storage	Availability	Triggered when the host that is part of a vSAN cluster is in disconnected state or not responding and vCenter Server does not know its state.
vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster.	Storage	Configuration	Triggered when vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster.

Table 8-164. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service.	Storage	Configuration	Triggered when vSAN cluster has issues in electing stats controller of vSAN Performance service.
vSAN cluster has multiple network partitions.	Network	Configuration	Triggered when vSAN cluster has multiple network partitions due to a network issue.
vSAN Cluster has multiple Stats DB objects which are creating conflicts and affecting vSAN Performance Service.	Storage	Configuration	Triggered when vSAN cluster has issues in electing stats controller of vSAN Performance service. This affects the functionality of vSAN Performance service.
vSAN disk group has incorrect deduplication and compression configuration.	Storage	Configuration	Triggered when vSAN disk group has incorrect deduplication and compression configuration.
vSAN has encountered an issue while reading the metadata of a physical disk.	Storage	Availability	Triggered when vSAN has encountered an issue while reading the metadata of a physical disk and cannot use this disk.
vSAN health service is not installed on the host.	Storage	Configuration	Triggered when vSAN health service is not installed on the host.
vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster.	Storage	Configuration	Triggered when vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster.
vSAN is unable to retrieve the physical disk information from host.	Storage	Availability	Triggered when vSAN is unable to retrieve the physical disk information from host. vSAN Health Service may not be working properly on this host.
vSAN Performance Service is not enabled.	Storage	Configuration	Triggered when vSAN Performance Service is not enabled.
vSAN Performance Service is unable to communicate and retrieve statistics from host.	Storage	Configuration	Triggered when vSAN Performance Service is unable to communicate and retrieve statistics from host.
vSAN Performance Service network diagnostic mode is enabled for more than 24 hours.	Storage	Configuration	Triggered when the network diagnostic mode in vSAN Performance Service is enabled for more than 24 hours.
vSAN Stretched cluster contains a witness host without a valid disk group.	Storage	Configuration	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.

Table 8-164. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN Stretched cluster does not contain a valid witness host.	Storage	Configuration	Triggered when vSAN Stretched cluster does not contain a valid witness host. This affects the operations of vSAN Stretched cluster.
vSAN Stretched cluster does not contain two valid fault domains.	Storage	Configuration	Triggered when vSAN Stretched cluster does not contain two valid fault domains.
vSAN Stretched cluster has inconsistent configuration for Unicast agent.	Storage	Configuration	Triggered when vSAN Stretched cluster contains multiple unicast agents. This means multiple unicast agents were set on non-witness hosts.
vSAN witness host has an invalid preferred fault domain.	Storage	Configuration	Triggered when vSAN witness host has an invalid preferred fault domain.
Witness host is a part of vSAN Stretched cluster.	Storage	Configuration	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.	Storage	Configuration	Triggered when witness host resides in one of the data fault domains. This affects the operations of vSAN Stretched cluster.
Witness appliance upgrade to vSphere 7.0 or higher with caution.	Storage	Configuration	Triggered when you want to upgrade the witness appliance to vSphere 7.0 or higher.
vSAN Support Insight is not enabled for the environment.	Storage	Configuration	Triggered when vSAN Support Insight is not enabled for the environment.
LSI 3108 controller's advanced configuration values is different from recommended values.	Storage	Configuration	Triggered when the LSI-3108 based controller configuration values differs from vSAN configuration recommended values.
vSAN Cluster Overall Health is Red.	Application	Performance	Triggered when the overall health of the vSAN cluster is impacted.
vSAN Cluster flash read cache reservation is approaching capacity.	Application	Performance	Triggered when the flash read cache reservation in a vSAN cluster is less than 20%. Cleared by adding more flash storage to the read-cache.
Some vSAN hosts are not compliant with the hyperconverged cluster configuration.	Storage	Configuration	Triggered when one of the host in vSAN cluster is not compliant with the hyperconverged cluster configuration.
Some vSAN hosts are not compliant for VMware vSphere Distributed Switch configuration.	Storage	Configuration	Triggered when one of the host in vSAN cluster is not compliant with the VMware vSphere Distributed Switch configuration.
Dual encryption is applied on virtual machines of a vSAN cluster.	Storage	Availability	Triggered when dual encryption is applied on a virtual machines of a vSAN cluster.

Table 8-165. vSAN Cluster Object Risk Alert Definitions

Alert	Alert Type	Alert Subtype	Description
After one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects	Storage	Capacity	Triggered when after one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects.
Capacity disk used for vSAN is smaller than 255 GB (default max component size).	Storage	Performance	Triggered when a capacity disk used for vSAN is smaller than 255 GB (default max component size), so virtual machines that run on the vSAN Datastore might experience disk space issues.
Capacity disk used for vSAN is smaller than 255 GB (default max component size).	Storage	Availability	Triggered when a capacity disk used for vSAN is smaller than 255 GB (default max component size), so virtual machines that run on the vSAN Datastore might experience disk space issues.
Controller with pass-through and RAID disks has issues.	Storage	Configuration	Triggered when a controller with pass-through and RAID disks has issues.
Disk format version of one or more vSAN disks is out of date	Storage	Configuration	Triggered when the 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.
ESXi host issues retrieving hardware info.	Storage	Configuration	Triggered when the ESXi host issues retrieving hardware info.
Firmware provider hasn't all its dependencies met or is not functioning as expected.	Storage	Configuration	Triggered when a firmware provider has not met all its dependencies or is not functioning as expected.
Host with inconsistent extended configurations is detected.	Storage	Configuration	Triggered when a host with inconsistent extended configurations is detected. vSAN cluster extended configurations are set as object repair timer is 60 minutes, site read locality is Enabled, customized swap object is Enabled, large scale cluster support is Disabled; For host with inconsistent extended configurations, vSAN cluster remediation is recommended, for host doesn't support any extended configuration, ESXi software upgrade is needed; And to make cluster scalability configuration take effect, host reboot could be required.
Inconsistent configuration (like dedup/compression, encryption) setup on hosts or disks with the cluster.	Storage	Configuration	Triggered when there is inconsistent configuration (like dedup/compression, encryption) setup on hosts or disks with the cluster.
Network adapter driver is not VMware certified.	Storage	Configuration	Triggered when the network adapter driver is not VMware certified.

Table 8-165. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
Network adapter firmware is not VMware certified.	Storage	Configuration	Triggered when the network adapter firmware is not VMware certified.
Network adapter is not VMware certified.	Storage	Configuration	Triggered when the network adapter is not VMware certified.
Network configuration of the vSAN iSCSI target service is not valid.	Storage	Availability	Triggered when the network configuration of the vSAN iSCSI target service is not valid. This health check validates the presence of the default vmknix for the vSAN iSCSI target service, and verifies that all the existing targets have valid vmknix configurations.
Non-vSAN disks are used for VMFS or Raw Device Mappings(RDMs).	Storage	Availability	Triggered when non-vSAN disks are used for VMFS or Raw Device Mappings (RDMs).
Number of vSAN components on a disk is reaching or has reached its limit.	Storage	Capacity	Triggered when the 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.
Number of vSAN components on a host is reaching or has reached its limit.	Storage	Capacity	Triggered when the 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.
One or more ESXi hosts in the cluster do not support CPU AES-NI or have it disabled.	Storage	Availability	Triggered when one or more hosts in the cluster do not support CPU AES-NI or have it disabled. As a result, the system might use the software encryption that is significantly slower than AES-NI.
RAID controller configuration has issues.	Storage	Configuration	Triggered when the RAID controller configuration has issues.
Storage I/O controller driver is not VMware certified	Storage	Configuration	Triggered when the stability and integrity of vSAN may be at risk as the storage I/O controller driver is not VMware certified.
Storage I/O controller drivers is not supported with the current version of ESXi running on the host	Storage	Configuration	Triggered when the 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.
Storage I/O Controller firmware not is VMware certified.	Storage	Configuration	Triggered when the storage I/O Controller firmware not is VMware certified.
Storage I/O controller is not compatible with the VMware Compatibility Guide	Storage	Configuration	Triggered when the 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.

Table 8-165. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
The current status of the Customer Experience Improvement Program (CEIP) not is enabled.	Storage	Availability	Triggered when the current status of the Customer Experience Improvement Program (CEIP) not is enabled.
The Internet connectivity is not available for vCenter Server.	Storage	Availability	Triggered when internet connectivity is not available for vCenter Server.
The resync operations are throttled on any hosts.	Storage	Configuration	Triggered when resync operations are throttled. Please clear the limit, unless you need it for particular cases like a potential cluster meltdown.
Time of hosts and VC are not synchronized within 1 minute.	Storage	Configuration	Triggered when the time of hosts and VC are not synchronized within 1 minute. Any difference larger than 60 seconds will lead this check to fail. If the check fails, it is recommended that you check the NTP server configuration.
vCenter Server or any of the ESXi hosts experience problems when connecting to Key Management Servers (KMS).	Storage	Availability	Triggered when the vCenter Server or any of the hosts experience problems when connecting to KMS.
vCenter server state was not pushed to ESXi due to vCenter server being out of sync.	Storage	Configuration	Triggered when the vCenter server state was not pushed to ESXi due to vCenter server being out of sync. During normal operation, the vCenter server state is regarded as source of truth, and ESXi hosts are automatically updated with the latest host membership list. When vCenter server is replaced or recovered from backup, the host membership list in vCenter server may be out of sync. This health check detects such cases, and alerts if vCenter server state was not pushed to ESXi due to vCenter server being out of sync. In such cases, first fully restore the membership list in vCenter server, and then perform 'Update ESXi configuration' action if required.
vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver.	Storage	Configuration	Triggered when the vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver.
vSAN build recommendation based on the available releases and VCG compatibility guide.	Storage	Availability	Triggered when the vSAN build is not compatible with available releases and VCG compatibility guide. This is the ESXi build that vSAN recommends as the most appropriate, given the hardware, its compatibility per the VMware Compatibility Guide and the available releases from VMware.

Table 8-165. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN build recommendation engine has all its dependencies met and is functioning as expected.	Storage	Availability	Triggered when the vSAN build recommendation engine has issues. The vSAN Build Recommendation Engine relies on the VMware compatibility guide and VMware release metadata for its recommendation. To provide build recommendations, it also requires VMware Update Manager service availability, internet connectivity, and valid credentials for my.vmware.com. This health check ensures that all dependencies are met and that the recommendation engine is functioning correctly.
vSAN Cluster disk space capacity is less than 5%	Storage	Capacity	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 disk space usage is approaching capacity	Storage	Capacity	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 is reaching or has reached its limit for components, free disk space and read cache reservations.	Storage	Capacity	Triggered when the vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations.
vSAN Cluster virtual disk count capacity is less than 5%.	Storage	Capacity	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 Cluster virtual disk count is approaching capacity.	Storage	Capacity	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 configuration for LSI 3108-based controller has issues.	Storage	Configuration	Triggered when the vSAN configuration for LSI 3108-based controller has issues.
vSAN disk group type (All-Flash or Hybrid) for the used SCSI controller is not VMware certified.	Storage	Configuration	Triggered when the vSAN disk group type (All-Flash or Hybrid) for the used SCSI controller is not VMware certified.
vSAN enabled hosts have inconsistent values for advanced configuration options.	Storage	Configuration	Triggered when some advanced configuration settings have different values on different hosts in the vSAN cluster.
vSAN firmware version recommendation based on the VCG.	Storage	Configuration	Triggered when the vSAN firmware version recommendation based on the VCG check has issues.

Table 8-165. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk.	Storage	Availability	Triggered when the vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk.
vSAN HCL DB auto updater is not working properly.	Storage	Configuration	Triggered when the vSAN HCL DB auto updater is not working properly. This means that vSAN cannot download and update its HCL DB automatically.
vSAN HCL DB is not up-to-date.	Storage	Configuration	Triggered when the vSAN HCL DB is not up-to-date.
vSAN Health Service is not able to find the appropriate controller utility for the storage controller on the ESXi host.	Storage	Availability	Triggered when the vSAN Health Service is not able to find the appropriate controller utility for the storage controller on the ESXi host.
vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.	Storage	Performance	Triggered when the 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 running low on the vital memory pool (slabs) needed for the operation of physical disks.	Storage	Performance	Triggered when the 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 using a physical disk which has high congestion value.	Storage	Performance	Triggered when the 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.
vSAN iSCSI target service home object has issues.	Storage	Availability	Triggered when the vSAN iSCSI target service home object has issues. This health check verifies the integrity of the vSAN iSCSI target service home object. It also verifies that the configuration of the home object is valid.

Table 8-165. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN iSCSI target service is not running properly or is not correctly enabled on the host.	Storage	Availability	Triggered when the vSAN iSCSI target service is not running properly or is not correctly enabled on the host. This health check verifies the service runtime status of the vSAN iSCSI target service, and checks whether the service is correctly enabled on each host.
vSAN performance service statistics database object is reporting issues.	Storage	Availability	Triggered when the vSAN performance service statistics database object is reporting issues.
vSphere cluster members do not match vSAN cluster members.	Storage	Configuration	Triggered when the vSphere cluster members do not match vSAN cluster members.

Table 8-166. vSAN Cluster Object Efficiency Alert Definitions

Alert	Alert Type	Alert Subtype	Description
vSAN Cluster flash read cache is approaching capacity.	Storage	Capacity	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%.	Storage	Capacity	Triggered when the Read Cache (RC) in the vSAN cluster reaches 95% of capacity. Cleared by adding flash storage to the read cache.

vSAN Adapter Instance Object Alert Definitions

Alerts on the vSAN Adapter Instance Object have health impact.

Alert	Alert Type	Alert Subtype	Description
vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.	Storage	Configuration	Triggered when the vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.

vSAN Disk Group Object Alert Definitions

Alerts on the vSAN Disk Group Object have efficiency impact.

Alert	Alert Type	Alert Subtype	Description
vSAN Disk Group read cache hit rate is less than 90%.	Storage	Performance	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%.	Storage	Capacity	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.

vSAN Host Object Alert Definitions

Alerts on the vSAN Host Object have security impact.

Alert	Alert Type	Alert Subtype	Description
vSAN host has encryption disabled, while the vSAN cluster has encryption enabled.	Storage	Configuration	Triggered when the vSAN host has encryption disabled, while the vSAN cluster has encryption enabled. Cleared by enabling encryption on vSAN host.
vSAN host encryption is enabled, while the vSAN cluster encryption is disabled.	Storage	Configuration	Triggered when the vSAN host has encryption enabled, while the vSAN cluster has encryption disabled. Cleared by enabling encryption on vSAN cluster.

vSAN Capacity Disk Object Alert Definitions

Alerts on the vSAN Capacity Disk object have security impact.

Alert	Alert Type	Alert Subtype	Description
vSAN capacity disk has encryption disabled, while the vSAN cluster has encryption enabled.	Storage	Configuration	Triggered when the vSAN capacity disk has encryption disabled, while the vSAN cluster has encryption enabled. Cleared by enabling encryption on vSAN capacity disk.
vSAN capacity disk encryption is enabled, while the vSAN cluster encryption is disabled.	Storage	Configuration	Triggered when the vSAN capacity disk has encryption enabled, while the vSAN cluster has encryption disabled. Cleared by enabling encryption on vSAN cluster.

Alert	Alert Type	Alert Subtype	Description
The free read cache reservations across the entire vSAN cluster are beyond the thresholds.	Storage	Capacity	Triggered when the flash read cache is exhausted. Note Flash read cache is only relevant to hybrid configurations and is not relevant on all-flash configurations.
Deployment of new virtual machines fails due to insufficient disk capacity	Storage	Capacity	Triggered when the disk capacity of the vSAN cluster exceeds the threshold value.

vSAN Cache Disk Object Alert Definitions

Alerts on the vSAN Cache Disk object have security impact.

Alert	Alert Type	Alert Subtype	Description
vSAN cache disk has encryption disabled, while the vSAN cluster has encryption enabled.	Storage	Configuration	Triggered when the vSAN cache disk has encryption disabled, while the vSAN cluster has encryption enabled. Cleared by enabling encryption on vSAN cache disk.
vSAN cache disk encryption is enabled, while the vSAN cluster encryption is disabled.	Storage	Configuration	Triggered when the vSAN cache disk has encryption enabled, while the vSAN cluster has encryption disabled. Cleared by enabling encryption on vSAN cluster.

vSAN File Service Alert Definitions

Alert	Alert Type	Alert Subtype	Description
vSAN File Service infrastructure health has issues.	Storage	Configuration	Triggered when there is an issue with file service infrastructure health state of an ESXi host in the vSAN cluster.
vSAN File Share health is not in a good state.	Storage	Configuration	Triggered when the vSAN File Share health is not in a good state.
Network File System (NFS) daemon is not running.	Storage	Configuration	Triggered when the NFS daemon process is not running.
Root File System is inaccessible.	Storage	Configuration	Triggered when the root file system does not repond to the file server.
File Server IP address not assigned.	Storage	Configuration	Triggered when IP address is not assigned to the file server.
vSAN File Server health is not in a good state.	Storage	Configuration	Triggered when the vSAN File Server health is not in a good state.

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 a link down state.	Symptoms include all of the following: <ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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 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 memory workload.	Symptoms include all of the following: <ul style="list-style-type: none"> ■ Virtual machine memory workload is at Warning/Immediate/Critical level ■ Anomaly is starting to/moderately/critically high 	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ ! 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] 	<ol style="list-style-type: none"> 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 disk I/O read latency problem.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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 	<ol style="list-style-type: none"> 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 Use vSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS.
Virtual machine has disk I/O write latency problem.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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) 	<ol style="list-style-type: none"> 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.	<p>Symptoms include all of the following:</p> <ul style="list-style-type: none"> ■ 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] 	<ol style="list-style-type: none"> 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
Not enough resources for vSphere HA to start the virtual machine.	Not enough resources for vSphere HA to start VM (Fault symptom).	<ol style="list-style-type: none"> 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)	<ol style="list-style-type: none"> 1 If the error information reports that a file is locked, the virtual machine might be powered on a host that the vSphere HA primary 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.
One or more virtual machine guest file systems are running out of disk space.	Symptoms include all of the following: <ul style="list-style-type: none"> ■ Guest file system usage at warning level ■ Guest file system usage at critical level 	Add a new virtual hard disk or expand the existing disk of the virtual machine. Before expanding the existing disk, remove all the snapshots. Once done, use a guest OS specific procedure to expand the file system on the new or expanded disk.
Virtual machine has CPU contention due to memory page swapping in the host.	Symptoms include all of the following: <ul style="list-style-type: none"> ■ Virtual machine CPU swap wait is at Critical level ■ Virtual machine CPU swap wait is at Immediate level ■ Virtual machine CPU swap wait is at Warning level 	<ol style="list-style-type: none"> 1 Set memory reservations for the virtual machine to prevent its memory from being swapped. 2 Verify that VMware Tools is installed and running, and that the balloon driver is enabled in the guest. Memory ballooning helps the host reclaim unused memory from the guest more effectively, and might avoid swapping. 3 Use vMotion to migrate this virtual machine to a different host or cluster.

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: <ul style="list-style-type: none">■ 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: <ul style="list-style-type: none"> ■ 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 is violating vSphere 5.5 hardening guide.	<ul style="list-style-type: none"> ■ 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 ■ Shrink virtual disk not disabled - diskShrink OR 	Fix the vSphere 5.5 hardening guide rule violations according to the recommendations in the vSphere Hardening Guide (XLSX).

Alert Definition	Symptom	Recommendations
	<ul style="list-style-type: none"> ■ 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 ■ Feature not exposed in vSphere: Unity-Windowcontents is not disabled OR 	

Alert Definition	Symptom	Recommendations
	<ul style="list-style-type: none"> Feature not exposed in vSphere: Unity-Unityactive is not disabled 	
Virtual machine has CPU contention due to multi-vCPU scheduling issues (co-stop) caused by snapshots	Symptoms include all of the following: <ul style="list-style-type: none"> Virtual machine CPU co-stop is at Warning level OR Virtual machine CPU co-stop is at Immediate level OR Virtual machine CPU co-stop is at Critical level And <ul style="list-style-type: none"> Virtual machine is powered off OR Virtual machine has at least one snapshot 	None.

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 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 the vCenter Server.	Duplicate object name found in the vCenter Server.	Ensure that 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.
VASA Provider(s) disconnected	One or more VASA Providers disconnected from vCenter.	If the VASA provider is inaccessible from the vCenter and you are getting an invalid certificate error then, see KB article: 2079087 . Contact the hardware vendor for further support.
Certificate for VASA Provider(s) will expire soon	One or more VASA Providers' certificates expire soon.	Contact the hardware vendor for getting support on the CA certificates and CRLs for VASA provider.
Refreshing CA certificates and CRLs for VASA Provider(s) failed	Refreshing CA certificates and CRLs for one or more VASA Providers failed.	Refresh the storage provider certificate as per the following document: <i>Refresh Storage Provider Certificates</i> . Contact the hardware vendor for further support. Note The <i>Refresh Storage Provider Certificates</i> is in the vSphere Storage 6.5 guide.
Virtual machine has memory contention caused by swap wait and high disk read latency.	Virtual Machine has a memory contention due to swap wait and high disk read latency.	Add more memory for the virtual machine and ensure that VMware Tools is running in the virtual machine.

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 due to multi-vCPU scheduling issues (co-stop) caused by too many vCPUs.	Virtual Machine experiences a high co-stop. The co-stop is the amount of time taken when the virtual machine is ready to run but is experiencing delay because of the co-vCPU scheduling contention. High co-stop occurs when too many vCPUs are configured for the virtual machine, and not enough physical CPUs are available to manage the co-vCPU scheduling.	Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended.

Datastore Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

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).	<p>The storage device path, for example, <code>vmhba35:C1:T0:L7</code>, contains several potential failure points: Path Element Failure Point</p> <p>-----</p> <p>vmhba35 HBA (Host Bus Adapter) C1 Channel T0 Target (storage processor port) L7 LUN (Logical Unit Number or Disk Unit).</p> <p>To determine the cause of the failure or to eliminate possible problems: Identify the available storage paths to the reported storage device by running <code>esxcfg-mpath - l</code>. For more information, see http://kb.vmware.com/kb/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 http://kb.vmware.com/kb/1003988. Determine whether the connectivity issue is with the iSCSI storage or the fiber storage.</p> <p>Troubleshoot the connectivity to the iSCSI storage by using the software initiator:</p> <ol style="list-style-type: none"> 1 Check whether a ping to the storage array fails from ESX. For more information, see http://kb.vmware.com/kb/1003486 2 Check whether a vmkping to each network portal of the storage array fails. For more information, see http://kb.vmware.com/kb/10037828. 3 Check that the initiator is registered on the array. For more information, contact your storage vendor.

Alert Definition	Symptom	Recommendations
		<p>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.</p> <p>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 http://kb.vmware.com/kb/1002301.</p> <p>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.</p> <p>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.</p>

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).	<p>The storage device path, for example, vmhba35:C1:T0:L7, contains several potential failure points:</p> <p>Path Element Failure Point</p> <p>-----</p> <p>vmhba35 HBA (Host Bus Adapter) C1 Channel T0 Target (storage processor port) L7 LUN (Logical Unit Number or Disk Unit).</p> <p>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 <code>esxcfg-mpath - l</code>. For more information, see http://kb.vmware.com/kb/1003973.</p> <p>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 http://kb.vmware.com/kb/1003988.</p> <p>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:</p> <ol style="list-style-type: none"> 1 Check whether a ping to the storage array fails from ESX. For more information, see http://kb.vmware.com/kb/1003486. 2 Check whether a vmkping to each network portal of the storage array fails. For more information, see http://kb.vmware.com/kb/10037828. 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
		<p>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 http://kb.vmware.com/kb/1002301.</p> <p>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.</p>

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: <ul style="list-style-type: none"> ■ Datastore space usage reaching warning/immediate/critical level ■ ! Datastore space growth above DT ■ Datastore space time remaining is low 	<ol style="list-style-type: none"> 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.

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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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.

vSphere Pod Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vSphere Pod objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk/Health

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Not enough resources for vSphere HA to start the Pod	Not enough resources for vSphere HA to start Pod	
One or more Pod guest file systems are running out of disk space	Symptom set is met when any of the symptoms are true: <ul style="list-style-type: none"> ■ Guest file system space usage at warning level ■ Guest file system space usage at critical level 	
Pod CPU usage is at 100% for an extended period of time	Pod sustained CPU usage is 100%	
Pod disk I/O read latency is high	Symptom set is met when any of the symptoms are true: <ul style="list-style-type: none"> ■ Pod disk read latency at Warning level ■ Pod disk read latency at Immediate level ■ Pod disk read latency at Critical level 	
Pod disk I/O write latency is high	Symptom set is met when any of the symptoms are true: <ul style="list-style-type: none"> ■ Pod disk write latency at Warning level ■ Pod disk write latency at Immediate level ■ Pod disk write latency at Critical level 	
Pod has CPU contention due to long wait for I/O events	Symptom set is met when any of the symptoms are true: <ul style="list-style-type: none"> ■ Pod CPU I/O wait is at Critical level ■ Pod CPU I/O wait is at Immediate level ■ Pod CPU I/O wait is at Warning level 	
Pod has CPU contention due to memory page swapping in the host	Symptom set is met when any of the symptoms are true. <ul style="list-style-type: none"> ■ Pod CPU swap wait is at Critical level ■ Pod CPU swap wait is at Immediate level ■ Pod CPU swap wait is at Warning level 	

Alert Definition	Symptoms	Recommendations
Pod has CPU contention due to multi-vCPU scheduling issues (co-stop) caused by too many vCPUs	<p>Alert is triggered when all of the symptom sets are true.</p> <ul style="list-style-type: none"> ■ Pod is powered off <p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU co-stop is at Critical level ■ Pod CPU co-stop is at Immediate level ■ Pod CPU co-stop is at Warning level 	
Pod has memory contention caused by swap wait and high disk read latency	<p>Alert is triggered when all of the symptom sets are true.</p> <p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU swap wait is at Warning level ■ Pod CPU swap wait is at Immediate level ■ Pod CPU swap wait is at Critical level <p>Symptom set is met when all of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod disk read latency at Warning level ■ VMware Tools is running ■ Pod does not have memory ballooning 	
Pod has memory contention due to memory compression, ballooning, or swapping	<p>Alert is triggered when all of the symptom sets are true:</p> <ul style="list-style-type: none"> ■ Pod memory limit is set <p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod memory contention is at Critical level ■ Pod memory contention is at Immediate level ■ Pod memory contention is at warning level ■ Pod memory is compressed ■ Pod memory ballooning is at Warning level ■ Pod memory ballooning is at Immediate level ■ Pod memory ballooning is at Critical level ■ Pod is using swap 	

Alert Definition	Symptoms	Recommendations
Pod is demanding more CPU than the configured limit	<p>Symptom set is met when all of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU limit is set ■ CPU Demand is greater than configured limit 	
Pod is experiencing memory compression, ballooning, or swapping due to memory limit	<p>Alert is triggered when all of the symptom sets are true.</p> <ul style="list-style-type: none"> ■ Pod memory limit is set ■ Pod memory demand exceeds configured memory limit <p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod memory is compressed ■ Pod memory ballooning is at Warning level ■ Pod memory ballooning is at Immediate level ■ Pod memory ballooning is at Critical level ■ Pod is using swap 	
Pod is in an invalid or orphaned state	<p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod is in invalid state ■ Pod is orphaned 	
Pod on a host with BIOS power management not set to OS controlled is facing CPU contention	<p>Alert is triggered when all of the symptom sets are true:</p> <ul style="list-style-type: none"> ■ Pod CPU contention at critical level <p>Symptom set is true when all of parent host system exhibit the following symptom.</p> <ul style="list-style-type: none"> ■ Host power management technology is not set to OS Controlled 	
Pod on a host with BIOS power management not set to OS controlled is facing CPU contention	<p>Alert is triggered when all of the symptom sets are true.</p> <p>Symptom set is met when all of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU contention is elevated ■ Pod CPU contention at critical level <p>Symptom set is true when all of parent host system exhibit the following symptom.</p> <ul style="list-style-type: none"> ■ Host power management technology is not set to OS Controlled 	

Alert Definition	Symptoms	Recommendations
Pod on a host with BIOS power management set to OS controlled is facing CPU contention	<p>Alert is triggered when all of the symptom sets are true.</p> <p>Symptom set is met when all of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU contention is elevated ■ Pod CPU contention is elevated <p>Symptom set is true when all of parent host system exhibit the following symptom.</p> <ul style="list-style-type: none"> ■ Host power management technology is not set to OS Controlled 	
Pod on a host with BIOS power management set to OS controlled is facing CPU contention	<p>Alert is triggered when all of the symptom sets are true.</p> <p>Symptom set is met when any of the symptoms are true.</p> <ul style="list-style-type: none"> ■ Pod CPU contention is elevated ■ Pod CPU contention is elevated ■ Pod CPU contention at critical level <p>Symptom set is true when all of parent host system exhibit the following symptom.</p> <ul style="list-style-type: none"> ■ Host power management technology is not set to OS Controlled 	
vSphere HA failed to restart a network isolated Pod	vSphere HA failed to restart a network isolated Pod	

VMware Cloud on AWS 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 **VMware Cloud on AWS** objects.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact

Risk

Criticality

Symptom-based

Alert Definition	Symptoms	Recommendations
Number of SDDCs in this organization is exceeding the supported configuration maximum	VMC Configuration Maximum limits breached. The number of SDDCs in this organization is over the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings. If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.
Number of hosts per SDDC is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The number of hosts in this SDDC is at the supported limit.	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of clusters per SDDC soft limit is exceeding the supported configuration maximum	VMC Configuration Maximum limits are maxed out. Maximum number of clusters soft limit is over the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMware Cloud on AWS Configuration Maximum guide. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings. If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.
Number of virtual machines per SDDC is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. Number of virtual machines per SDDC is at the supported maximum	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of linked VPCs in this SDDC is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The number of linked VPCs in this SDDC is at the supported limit.	Please refer to VMC on AWS guide listed here .
Number of SDDCs in this organization is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The number of SDDCs in this organization is at the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings. If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.

Alert Definition	Symptoms	Recommendations
Number of Public IP Addresses (Elastic IPs) per organization is exceeding the supported configuration maximum	VMC Configuration Maximum limits breached. The Number of Public IP Addresses (Elastic IPs) per organization is over the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.
Number of clusters per SDDC hard limit is at supported configuration maximum	VMC Configuration Maximum limits are maxed out. Maximum number of clusters hard limit is at supported configuration maximum	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of virtual machines per SDDC is exceeding the supported configuration maximum	VMC Configuration Maximum limits are breached. The number of virtual machines per SDDC is exceeding the supported maximum	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of linked VPCs in this SDDC is exceeding the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The number of linked VPCs in this SDDC is over the supported limit.	Please refer to VMC on AWS guide listed here .
Number of clusters per SDDC hard limit is exceeding the supported configuration maximum	VMC Configuration Maximum limits are maxed out. Maximum number of clusters hard limit is over the supported limit.	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of clusters per SDDC soft limit is at supported configuration maximum	VMC Configuration Maximum limits are maxed out. Maximum number of clusters soft limit is at supported configuration maximum	<ul style="list-style-type: none"> ■ Please refer to VMware Cloud on AWS Configuration Maximum guide. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.

Alert Definition	Symptoms	Recommendations
Number of hosts per organization is exceeding the supported configuration maximum	VMC Configuration Maximum limits are breached. The number of hosts in this organization is over the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.
Number of hosts per organization is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The number of hosts in this organization is at the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.
Number of hosts per SDDC is exceeding the supported configuration maximum	VMC Configuration Maximum limits are breached. The number of hosts in this SDDC is over the supported limit.	Please refer to VMware Cloud on AWS Configuration Maximum guide.
Number of Public IP Addresses (Elastic IPs) per organization is at the supported configuration maximum	VMC Configuration Maximum limits are maxed out. The Number of Public IP Addresses (Elastic IPs) per organization is at the supported limit.	<ul style="list-style-type: none"> ■ Please refer to VMC on AWS guide listed here. ■ A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in vRealize Operations automatically, then refer to the KB article, KB 2059936.

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 Manager and 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 8-167. 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 8-168. 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 8-169. 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 the Adapter level.

Virtual Machine Properties

vRealize Operations Manager collects configuration, runtime, CPU, memory, network I/O, and properties about summary use for virtual machine objects. Properties are collected with the first cycle of data collection. Once collected, the next property collection occurs only when there is data change. In case of no data change, no property is collected.

Table 8-170. 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 8-171. 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 8-172. Properties Collected for Guest File Systems

Property Key	Property Name	Description
guestfilesystem capacity_property	Guest File System stats Guest File System Capacity Property	This property is disabled by default.
guestfilesystem capacity_property_total	Guest File System stats Total Guest File System Capacity Property(gb)	This property is disabled by default.

Table 8-173. Properties Collected for Disk Space Objects

Property Key	Property Name	Description
diskspace snapshot creator	Disk Space Snapshot Creator	This property is disabled by default.
diskspace snapshot description	Disk Space Snapshot Description	This property is disabled by default.

Table 8-174. Configuration Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
config name	Name	Name
config guestFullName	Guest OS from vCenter	This property is set by the vCenter during the VM creation. It may differ from the value of the Guest/
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)
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)

Table 8-174. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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)
config security disable_unexposed_features_memfs	Disable unexposed features - memfs (isolation.tools.memSchedFakeSampleStats.disable)	Disable unexposed features - memfs (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)

Table 8-174. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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)
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 (isolation.tools.guestDnDVersionSet.disable)	Disable unexposed features - versionset (isolation.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)

Table 8-174. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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
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
config faultTolerant	config faultTolerant	

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.

Table 8-175. Runtime Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
runtime memoryCap	Memory Capacity	Memory Capacity

Table 8-176. 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
cpuspeed	CPU	CPU Speed

Table 8-177. Memory Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
mem host_limit	VM Limit	Mem Machine Limit
mem host_reservation	Memory VM Reservation(kb)	This property is disabled by default.

Table 8-178. 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 vnic_label	Network:<ID> Label	This property is disabled by default.
net nvp_vm_uuid	Network I/O NVP VM UUID	This property is disabled by default.
net vnic_type	Network I/O Virtual NIC Type	This property is disabled by default.
net ipv6_address	Network IPv6 Address	This property is disabled by default.
net ipv6_prefix_length	Network IPv6 Prefix Length	This property is disabled by default.
net default_gateway	Network Network I/O Default Gateway	This property is disabled by default.
net subnet_mask	Network Subnet Mask	This property is disabled by default.

Table 8-179. 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

Table 8-179. Summary Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
summary parentDatacenter	Parent data center	Parent data center
summary parentVcenter	Parent vCenter	Parent vCenter
summary guest fullName	Guest OS Full Name	This property is provided by the VMware Tools. It will differ to the value set in vCenter if the Guest OS was upgraded, or if a different Guest OS was installed.
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
summary config appliance	summary config appliance	
summary config productName	Summary Configuration Product Name	

Table 8-180. Virtual Disk Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
virtualDisk configuredGB	Virtual Disk Configured(GB)	Virtual Disk configured disk space.
virtualDisk datastore	Virtual Disk Datastore	Datastore.
virtualDisk fileName	Virtual Disk File Name	This property is disabled by default.
virtualDisk label	Virtual Disk Label	Device label.

Table 8-181. Datastore Properties Collected for Virtual Machine Properties

Property Key	Property Name	Description
datastore maxObservedNumberRead	Datastore I/O Highest Observed Number of Read Requests	
datastore maxObservedNumberWrite	Datastore I/O Highest Observed Number of Write Requests	
datastore maxObservedOIO	Datastore I/O Highest Observed Outstanding Requests	
datastore maxObservedRead	Datastore I/O Highest Observed Read Rate(kbps)	
datastore maxObservedWrite	Datastore I/O Highest Observed Write Rate(kbps)	

Datastore properties collected for virtual machine objects have been disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

Host System Properties

vRealize Operations Manager collects configuration, hardware, runtime, CPU, network I/O, and properties about summary use for host system objects.

Table 8-182. 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

Table 8-182. Configuration Properties Collected for Host System Objects (continued)

Property Key	Property Name	Description
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)
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.
config security tlsdisabledprotocols	TLS Disabled Protocols	TLS Disabled Protocols

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.

Table 8-183. Cost Properties Collected for Host System Objects

Property Key	Property Name	Description
cost isServerLeased	Is Server Leased	Is Server Leased
cost remainingDepreciationMonths	Remaining Depreciation Months	Remaining number of Depreciation Months
cost serverPurchaseCost	Server Purchase Cost	Server Purchase Cost is displayed
cost serverPurchaseDate	Server Purchase Date	Server Purchase Date is displayed

Table 8-184. 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
hardware vendor	Hardware Vendor	Indicates the hardware manufacturer

Table 8-185. 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 8-186. 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 8-187. 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 8-188. 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 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

Table 8-189. 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 8-190. 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 8-191. Datastore Properties Collected for Host System Objects

Property Key	Property Name	Description
datastore maxObservedNumberRead	Datastore I/O Highest Observed Number of Read Requests	
datastore maxObservedNumberWrite	Datastore I/O Highest Observed Number of Write Requests	
datastore maxObservedOIO	Datastore I/O Highest Observed Outstanding Requests	
datastore maxObservedRead	Datastore I/O Highest Observed Read Rate(kbps)	
datastore maxObservedWrite	Datastore I/O Highest Observed Write Rate(kbps)	
net discoveryProtocol cdp timeToLive	Network I/O Discovery Protocol Cisco Discovery Protocol Time to Live	
net discoveryProtocol lldp timeToLive	Network I/O Discovery Protocol Link Layer Discovery Protocol Time to Live	

Datastore properties collected for host system objects have been disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

Table 8-192. Storage Path Properties Collected for Host System Objects

Property Key	Property Name	Description
storageAdapter port_WWN	Storage Adapter Port WWN	The port world wide name for storage adapter. Available for FC adapters only.

Cluster Compute Resource Properties

vRealize Operations Manager collects configuration and summary properties for cluster compute resource objects.

Table 8-193. Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
config name	Name	Name

Table 8-194. Summary Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
summary parentDatacenter	Parent data center	Parent data center
summary parentVcenter	Parent vCenter	Parent vCenter
summary customTag customTagValue	Value	Custom Tag Value
summary tag	vSphere Tag	vSphere Tag Name

Table 8-195. 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 Behavior	Default DRS Behavior
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 Behavior	Default DPM Behavior

Table 8-195. DR, DAS, and DPM Configuration Properties Collected for Cluster Compute Resource Objects (continued)

Property Key	Property Name	Description
configuration drsConfig pctIdleMBInMemDemand	Cluster Configuration DRS Configuration Idle Consumed Memory	
configuration drsConfig targetBalance	Cluster Configuration DRS Configuration Tolerable imbalance threshold	

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 8-196. 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 8-197. 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

Table 8-197. CPU Usage Properties Collected for Resource Pool Objects (continued)

Property Key	Property Name	Description
cpulshares	CPU Shares	CPU Shares
cpulcorecount_provisioned	Provisioned vCPU(s)	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.

Table 8-198. 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 8-199. 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 8-200. Configuration Properties Collected for Data Center Objects

Property Key	Property Name	Description
config name	Name	Name

Table 8-201. 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 8-202. 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 vmDiskAntiAffinityRules	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 8-203. Configuration Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
config name	Name	Name

Table 8-204. 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 8-205. Configuration Properties Collected for Distributed Virtual Port Group Objects

Property Key	Property Name	Description
config name	Name	Name
Configuration Uplink	Uplink	Indicates whether the portgroup is uplink portgroup.

Table 8-206. 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 8-207. Configuration Properties Collected for Datastore Objects

Property Key	Property Name	Description
config name	Name	Name

Table 8-208. 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
summary path	Summary Path	
summary scsiAdapterType	Summary SCSI Adapter Type	This property is disabled by default.
summary aliasOf	Summary Alias Of	Indicates whether the datastore is an alias of another. The published value is the container ID of the datastore for which it is an alias. Note This property may have 2 values. It's either "none", that means the datastore is not an alias of another datastore, or datastore <containerID> that is the Container ID of the datastore for which this is an alias.

Table 8-209. 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	Datastore I/O Highest Observed Number of Read Requests	Disabled
* datastore maxObservedNumberWrite	Datastore I/O Highest Observed Number of Write Requests	Disabled
* datastore maxObservedOIO	Datastore I/O Highest Observed Outstanding Requests	Disabled
* datastore maxObservedRead	Datastore I/O Highest Observed Read Latency	Disabled
* datastore maxObservedReadLatency	Datastore I/O Highest Observed Read Latency	Disabled

Table 8-209. Datastore Properties Collected for Datastore Objects (continued)

Property Key	Property Name	Description
* datastore maxObservedWrite	Datastore I/O Highest Observed Write Latency	Disabled
* datastore maxObservedWriteLatency	Datastore I/O Highest Observed Write Latency	Disabled

Table 8-210. Datastore Properties Collected for vVol Datastore Objects

Property Key	Property Name	Description
storageArray modelId	Storage Array Model	Storage array model of vVol datastore. Note This property is published for vVol datastores only and is available starting from vCenter version 6.0.
storageArray name	Storage Array Name	Storage array name of vVol datastore. Note This property is published for vVol datastores only and is available starting from vCenter version 6.0.
storageArray id	Storage Array ID	Storage array ID of vVol datastore. Note This property is published for vVol datastores only and is available starting from vCenter version 6.0.
storageArray vendorId	Storage Array Vendor	Storage array vendor of vVol datastore. Note This property is published for vVol datastores only and is available starting from vCenter version 6.0.
protocolEndpoints name	Protocol Endpoints Name	Protocol endpoint's name of vVol datastore. Note This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.0.

Table 8-210. Datastore Properties Collected for vVol Datastore Objects (continued)

Property Key	Property Name	Description
protocolEndpoints type	Protocol Endpoints Type	Protocol endpoint's type of vVol datastore. Note This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.5.
protocolEndpoints hosts	Protocol Endpoints Hosts	Hosts associated with protocol endpoint of vVol datastore. Note This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.0.

Datastore properties marked with an asterisk (*) have been disabled in this version of vRealize Operations Manager . This means that they do not collect data by default.

vSphere Pod Properties

vRealize Operations Manager collects summary and event properties for vSphere Pods.

Table 8-211. Summary Properties Collected for vSphere Pod Objects

Property Key	Localized Name	Description
config name	Configuration Name	Resource name.
config guestFullName	Configuration Guest OS from vCenter	This is the value provided by vCenter. vCenter set it during VM creation. The value may not match the value inside the Guest.
config version	Configuration Version	Virtual Machine Version.
config createDate	Configuration Creation Date	Object Creation Date.
config numVMDKs	Configuration Number of Virtual Disks	Number of Virtual Disks.
config faultTolerant	Configuration Fault Tolerant	Fault tolerance enabled.
config ft_role	Configuration FT Role	Role of the VM in Fault Tolerance Group.
config ft_peer_vm	Configuration FT Peer VM	Peer of the VM in Fault Tolerance Group.
config hardware numCpu	Configuration Hardware Number of virtual CPUs	Number of virtual CPUs.
config hardware memoryKB	Configuration Hardware Memory	Memory.
config hardware thinEnabled	Configuration Hardware Thin Provisioned Disk	Thin Provisioned Disk.

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config hardware numCoresPerSocket	Configuration Hardware Number of CPU cores per socket	Number of CPU cores per virtual socket.
config hardware numSockets	Configuration Hardware Number of virtual sockets	Number of virtual sockets.
config hardware diskSpace	Configuration Hardware Disk Space	Disk space metrics.
config cpuAllocation reservation	Configuration CPU Resource Allocation Reservation	N/A
config cpuAllocation limit	Configuration CPU Resource Allocation Limit	
config cpuAllocation shares shares	Configuration CPU Resource Allocation Shares Shares	
config memoryAllocation reservation	Configuration Memory Resource Allocation Reservation	
config memoryAllocation limit	Configuration Memory Resource Allocation Limit	
config memoryAllocation shares shares	Configuration Memory Resource Allocation Shares Shares	Memory Hot Add Configuration.
config extraConfig mem_hotadd	Configuration Extra Configuration Memory Hot Add	
config extraConfig vcpu_hotadd	Configuration Extra Configuration vCPU Hot Add	
config extraConfig vcpu_hotremove	Configuration Extra Configuration vCPU Hot Remove	
config extraConfig mem_tps_share	Configuration Extra Configuration VM MEM TPS	
config security disable_autoinstall	Configuration Security Disable tools auto install (isolation.tools.autoInstall.disable)	N/A
config security disable_console_copy	Configuration Security Disable console copy operations (isolation.tools.copy.disable)	
config security disable_console_dnd	Configuration Security Disable console drag and drop operations (isolation.tools.dnd.disable)	

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config security enable_console_gui_options	Configuration Security Enable console GUI operations (isolation.tools.setGUIOptions.enable)	
config security disable_console_paste	Configuration Security Disable console paste operations (isolation.tools.paste.disable)	
config security disable_disk_shrinking_shrink	Configuration Security Disable virtual disk shrink (isolation.tools.diskShrink.disable)	
config security disable_disk_shrinking_wiper	Configuration Security Disable virtual disk wiper (isolation.tools.diskWiper.disable)	
config security disable_hgfs	Configuration Security Disable HGFS file transfers (isolation.tools.hgfsServerSet.disable)	
config security disable_independent_nonpersistent	Configuration Security Avoid using independent nonpersistent disks (scsiX:Y.mode)	
config security enable_intervm_vmci	Configuration Security Enable VM-to-VM communication through VMCI (vmci0.unrestricted)	
config security enable_logging	Configuration Security Enable VM logging (logging)	
config security disable_monitor_control	Configuration Security Disable VM Monitor Control (isolation.monitor.control.disable)	
config security enable_non_essential_3D_features	Configuration Security Enable 3D features on Server and desktop virtual machines (mks.enable3d)	
config security disable_unexposed_features_autologon	Configuration Security Disable unexposed features - autologon (isolation.tools.ghi.autologon.disable)	

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config security disable_unexposed_features_biosbbs	Configuration Security Disable unexposed features - biosbbs (isolation.bios.bbs.disable)	
config security disable_unexposed_features_getcreds	Configuration Security Disable unexposed features - getcreds (isolation.tools.getCreds.disable)	
config security disable_unexposed_features_launchmenu	Configuration Security Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change)	
config security disable_unexposed_features_memsfss	Configuration Security Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable)	
config security disable_unexposed_features_protocolhandler	Configuration Security Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)	
config security disable_unexposed_features_shellaction	Configuration Security Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)	
config security disable_unexposed_features_toporequest	Configuration Security Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)	
config security disable_unexposed_features_trashfolderstate	Configuration Security Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)	
config security disable_unexposed_features_trayicon	Configuration Security Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable)	

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config security disable_unexposed_features_unity	Configuration Security Disable unexposed features - unity (isolation.tools.unity.disable)	
config security disable_unexposed_features_unity_in terlock	Configuration Security Disable unexposed features - unity-interlock (isolation.tools.unityInterlock Operation.disable)	
config security disable_unexposed_features_unity_ta skbar	Configuration Security Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar. disable)	
config security disable_unexposed_features_unity_u nityactive	Configuration Security Disable unexposed features - unity-unityactive (isolation.tools.unityActive.di sable)	
config security disable_unexposed_features_unity_wi ndowcontents	Configuration Security Disable unexposed features - unity-windowcontents (isolation.tools.unity.window Contents.disable)	
config security disable_unexposed_features_unitypus h	Configuration Security Disable unexposed features - unitypush (isolation.tools.unity.push.up date.disable)	
config security disable_unexposed_features_versiong et	Configuration Security Disable unexposed features - versionget (isolation.tools.vmxDnDVersi onGet.disable)	
config security disable_unexposed_features_versions et	Configuration Security Disable unexposed features - versionset (solation.tools.guestDnDVers ionSet.disable)	
config security disable_vix_messages	Configuration Security Disable VIX messages from the VM (isolation.tools.vixMessage.di sable)	

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config security enable_vga_only_mode	Configuration Security Disable all but VGA mode on virtual machines (svga.vgaOnly)	
config security limit_console_connection	Configuration Security Limit number of console connections (RemoteDisplay.maxConnection)	
config security limit_log_number	Configuration Security Limit number of log files (log.keepOld)	
config security limit_log_size	Configuration Security Limit log file size (log.rotateSize)	
config security limit_setinfo_size	Configuration Security Limit VMX file size (tools.setInfo.sizeLimit)	
config security enable_console_VNC	Configuration Security Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled)	
config security disable_device_interaction_connect	Configuration Security Disable unauthorized removal, connection of devices (isolation.device.connectable.disable)	
config security disable_device_interaction_edit	Configuration Security Disable unauthorized modification of devices (isolation.device.edit.disable)	
config security enable_host_info	Configuration Security Enable send host information to guests (tools.guestlib.enableHostInfo)	
config security network_filter_enable	Configuration Security Enable dvfilter network APIs (ethernetX.filterY.name)	
config security vm_safe_cpumem_agentaddress	Configuration Security VMsafe CPU/memory APIs - IP address (vm_safe.agentAddress)	

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
config security vmsafe_cpumem_agentport	Configuration Security VMsafe CPU/memory APIs - port number (vmsafe.agentPort)	
config security vmsafe_cpumem_enable	Configuration Security Enable VMsafe CPU/memory APIs (vmsafe.enable)	
config security disconnect_devices_floppy	Configuration Security Disconnect floppy drive	
config security disconnect_devices_cd	Configuration Security Disconnect CD-ROM	
config security disconnect_devices_usb	Configuration Security Disconnect USB controller	
config security disconnect_devices_parallel	Configuration Security Disconnect parallel port	
config security disconnect_devices_serial	Configuration Security Disconnect serial port	
config security pci_device_configured	Configuration Security DCUI timeout	
runtime memoryCap	Runtime Memory Capacity	Memory Capacity.
cpullimit	CPU CPU Limit	CPU Limit.
cpu reservation	CPU CPU reservation	CPU Reservation.
cpu speed	CPU CPU	CPU Speed.
mem host_reservation	Memory Host Active	Machine Active.
mem host_active	Memory Host Usage	Machine Usage.
net mac_address	Network Mac Address	N/A
net ip_address	Network IP Address	
net subnet_mask	Network Subnet Mask	
net ipv6_address	Network IPv6 Address	IPv6 Address.
net ipv6_prefix_length	Network IPv6 Prefix Length	IPv6 Prefix Length.
net default_gateway	Network Default Gateway	N/A
net nvp_vm_uuid	Network NVP VM UUID	
net vnic_type	Network Virtual NIC Type	Virtual Machine's network adapter type.
net vnic_label	Network Label	Device label.

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
summary UUID	Summary UUID	Instance UUID in vCenter that uniquely identify all virtual machine instances.
summary MOID	Summary MOID	Managed object ID in vCenter. This is unique in scope of vCenter.
summary swapOnlyDatastore	Summary Datastore with only swap file	Datastore containing only the swap file and no other files from this VM.
summary customTag customTagValue	Summary Custom Tag Value	Custom Tag Value.
summary tag	Summary vSphere Tag	vSphere Tag Name.
summary tag.Json	Summary vSphere Tag Json	vSphere Tag in Json format.
summary folder	Summary vSphere Folder	vSphere Folder Name.
summary parentCluster	Summary Parent Cluster	Parent Cluster.
summary parentHost	Summary Parent Host	Parent Host.
summary parentDatacenter	Summary Parent Datacenter	Parent Datacenter.
summary parentNamespace	Summary Parent Namespace	Parent Namespace.
summary parentVcenter	Summary Parent vCenter	Parent vCenter.
summary parentFolder	Summary Parent Folder	Parent Folder.
summary datastore	Summary Datastore(s)	Datastore(s).
summary guest fullName	Summary Guest Operating System Guest OS from Tools	This is the value provided by VMware Tools. This value will differ to the value set in vCenter if the Guest OS was upgraded, or a different Guest OS was installed.
summary guest ipAddress	Summary Guest Operating System Guest OS IP Address	Guest OS IP Address.
summary guest hostName	Summary Guest Operating System Hostname	Hostname of the guest operating system, if known.
summary guest toolsRunningStatus	Summary Guest Operating System Tools Running Status	Guest Tools Running Status.
summary guest toolsVersionStatus2	Summary Guest Operating System Tools Version Status	Guest Tools Version Status 2.
summary guest toolsVersion	Summary Guest Operating System Tools Version	VM tools version installed on guest OS.
summary guest vrealize_operations_agent_id	Summary Guest Operating System vRealize Operations Agent ID	An ID to identify a VM in Agent Adapter's world.
summary guest vrealize_operations_euc_agent_id	Summary Guest Operating System vRealize Operations Euc Agent ID	An ID to identify a VM in Agent Adapter's world.

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
summary config numEthernetCards	Summary Configuration Number of NICs	Number of NICs.
summary config productName	Summary Configuration Product Name	Product Name.
summary config appliance	Summary Configuration Appliance	Appliance.
summary runtime isIdle	Summary Runtime Idleness indicator	This property indicates whether the monitored instance is idle or not.
summary runtime powerState	Summary Runtime Power State	Power State.
summary runtime connectionState	Summary Runtime Connection State	Connection State.
guestfilesystem capacity_property	Guest File System Guest File System Capacity Property	Total capacity of guest file system as a property.
guestfilesystem capacity_property_total	Guest File System Total Capacity Property	Total capacity of guest file system as a property.
virtualDisk datastore	Virtual Disk Datastore	Datastore.
virtualDisk configuredGB	Virtual Disk Configured	Virtual Disk configured disk space.
virtualDisk label	Virtual Disk Label	Device Label.
virtualDisk fileName	Virtual Disk File Name	Virtual Disk file name.
diskspace snapshot mor	Disk Space Snapshot Managed Object Reference	Managed Object Reference.
diskspace snapshot name	Disk Space Snapshot Name	Snapshot name.
diskspace snapshot numberOfDays	Disk Space Snapshot Number of Days Old	Number of days since snapshot creation.
diskspace snapshot snapshotAge	Disk Space Snapshot Age (Days)	Virtual Machine's topmost snapshot age in days.
diskspace snapshot creator	Disk Space Snapshot Creator	Creator.
diskspace snapshot description	Disk Space Snapshot Description	Snapshot description.
vsan policy compliance	vSAN VM Storage Policies Compliance	Compliance status of the VM storage object.
datastore maxObservedNumberRead	Datastore Highest Observed Number of Read Requests	Highest Observed Number of Read Requests.
datastore maxObservedRead	Datastore Highest Observed Read Rate	Highest Observed Read Rate (KBps).
datastore maxObservedNumberWrite	Datastore Highest Observed Number of Write Requests	Highest Observed Number of Write Requests.

Table 8-211. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
datastore maxObservedWrite	Datastore Highest Observed Write Rate	Highest Observed Write Rate (KBps).
datastore maxObservedOIO	Datastore Highest Observed Outstanding Requests	Highest Observed Outstanding Requests.

Namespace Properties

vRealize Operations Manager collects summary and event properties for Namespace.

Table 8-212. Summary Properties Collected for Namespace Objects

Property Key	Localized Name	Description
config name	Configuration Name	Resource name
config resourceLimits namespace cpu	Configuration Resource Limits Namespaces CPU	CPU
config resourceLimits namespace mem	Configuration Resource Limits Namespaces Memory	Memory
config resourceLimits namespace diskspace	Configuration Resource Limits Namespaces Disk Space	Disk space metrics
config resourceLimits containers cpu_request	Configuration Resource Limits Containers CPU Request	CPU Request Default
config resourceLimits containers cpu_limit	Configuration Resource Limits Containers CPU Limit	CPU Limit Default
config resourceLimits containers mem_request	Configuration Resource Limits Containers Memory Request	Memory Request Default
config resourceLimits containers mem_limit	Configuration Resource Limits Containers Memory Limit	Memory Limit Default
config objectLimits compute pod_count	Configuration Object Limits Compute Pods	Number of Pods
config objectLimits compute deployment_count	Configuration Object Limits Compute Deployments	Deployments
config objectLimits compute job_count	Configuration Object Limits Compute Jobs	Jobs
config objectLimits compute daemon_sets	Configuration Object Limits Compute Daemon Sets	Daemon Sets
config objectLimits compute replica_sets	Configuration Object Limits Compute Replica Sets	Replica Sets

Table 8-212. Summary Properties Collected for Namespace Objects (continued)

Property Key	Localized Name	Description
config objectLimits compute replication_controllers	Configuration Object Limits Compute Replication Controllers	Replication Controllers
config objectLimits compute stateful_sets	Configuration Object Limits Compute Stateful Sets	Stateful Sets
config objectLimits storage config_maps	Configuration Object Limits Storage Config Maps	Config Maps
config objectLimits storage secret_count	Configuration Object Limits Storage Secrets	Secrets
config objectLimits storage persistent_volume_claim	Configuration Object Limits Storage Persistent Volume Claim	Persistent Volume Claim
config objectLimits network services	Configuration Object Limits Network Services	Services
summary parentDatacenter	Summary Parent Datacenter	Parent Datacenter
summary parentCluster	Summary Parent Cluster	Parent Cluster
summary parentVcenter	Summary Parent vCenter	Parent vCenter
mem limit	Memory Memory limit	Memory limit
mem reservation	Memory Memory reservation	Memory reservation
mem expandable_reservation	Memory Memory expandable reservation	Memory Expandable Reservation
mem shares	Memory Memory Shares	Memory Shares
cpu limit	CPU CPU Limit	CPU Limit
cpu reservation	CPU CPU Reservation	CPU Reservation
cpu expandable_reservation	CPU CPU expandable reservation	CPU expandable Reservation
cpu shares	CPU CPU Shares	CPU Shares
cpu corecount_provisioned	CPU Provisioned vCPU(s)	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.

Tanzu Kubernetes cluster Properties

vRealize Operations Manager collects summary and event properties for Tanzu Kubernetes clusters.

Table 8-213. Summary Properties Collected for Tanzu Kubernetes cluster Objects

Property Key	Localized Name	Description
config name	Configuration Name	Resource name
config cpuAllocation reservation	Configuration CPU Resource Allocation Reservation	N/A
config cpuAllocation limit	Configuration CPU Resource Allocation Limit	N/A
config cpuAllocation expandableReservation	Configuration CPU Resource Allocation Expandable Reservation	N/A
config cpuAllocation shares shares	Configuration CPU Resource Allocation Shares Shares	N/A
config memoryAllocation reservation	Configuration Memory Resource Allocation Reservation	N/A
config memoryAllocation limit	Configuration Memory Resource Allocation Limit	N/A
config memoryAllocation expandableReservation	Configuration Memory Resource Allocation Expandable Reservation	N/A
config memoryAllocation shares shares	Configuration Memory Resource Allocation Shares Shares	N/A
cpu limit	CPU CPU Limit	CPU Limit
cpu reservation	CPU CPU Reservation	CPU Reservation
cpu expandable_reservation	CPU CPU expandable reservation	CPU expandable Reservation
cpu shares	CPU CPU Shares	CPU Shares
cpu corecount_provisioned	CPU Provisioned vCPU(s)	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.
mem limit	Memory Memory limit	Memory limit
mem reservation	Memory Memory reservation	Memory reservation
mem expandable_reservation	Memory Memory expandable reservation	Memory Expandable Reservation
mem shares	Memory Memory Shares	Memory Shares
summary parentDatacenter	Summary Parent Datacenter	Parent Datacenter
summary parentNamespace	Summary Parent Namespace	Parent Namespace

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 8-214. 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 – Primary, 1 – Replica, 4 – Client..
ShardRedundancyLevel	Shard redundancy level	The target number of redundant copies for Object data.
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 8-215. 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 8-216. Memory Properties Collected for Node Objects

Property Key	Property Name	Description
mem RAM	System RAM	System RAM

Table 8-217. Service Properties Collected for Node Objects

Property Key	Property Name	Description
service proclpid	Process ID	Process ID

Remote Collector Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager remote collector objects.

Table 8-218. 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 8-219. Memory Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
mem RAM	System RAM	System RAM

Table 8-220. Service Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
service proclpid	Process ID	Process ID

Service Discovery Properties

vRealize Operations Manager displays object properties for service discovery.

Service Discovery Adapter Instance Properties

vRealize Operations Manager displays the following properties for the service discovery adapter instance.

Table 8-221. Service Discovery Adapter Instance Properties

Property Name	Description
Action Identifier	An FQDN and IP pair of the end point vCenter Server that is used to identify the adapter instance that has to run actions on the vCenter Server.
Included Services	A list of user-defined services. The list entries are (service name, port, display name) triples separated by a new line.

Virtual Machine Properties

vRealize Operations Manager displays the following properties for virtual machines.

Table 8-222. Virtual Machine Properties

Property Name	Description
Guest OS Services Authentication Method	Refers to the VM guest operating system authentication method. The guest operating system can be authenticated either via a common user/password or a guest alias.
Guest OS Services Discovery Status	Reflects the result of service discovery operation on the VM's guest operating system.
Guest OS Services Authentication Status	Guest operating system authentication status.
Guest OS Services Inbound Ports	List of VM inbound ports. These are the ports on which the discovered services are listening.
SRM Info Protection Group	Protection group to which the VM belongs.
SRM Info Recovery Plans	List of recovery plans covering the VM.

Services Properties

vRealize Operations Manager displays the following properties for services.

Table 8-223. Services Properties

Property Name	Description
Type	The name of the service type.
Install Path	The install path.
Ports	List of service listening ports.
Virtual Machine	The name of the parent VM.
Virtual Machine MOID	The MOID of the VM.
Version	Version of the discovered service.
Is Application Member	Indicates that the service is a member of the group of services forming an application.
Category	Category of the service.
Process Name	Name of the process.
Connection Type	If there is a remote process that was connected to one of the listening ports of the given service, then the property's value is set to <i>Incoming</i> . If not, it is set to <i>Outgoing</i> . If there is no connection to another service, then the value of the property is set to <i>N/A</i> .
Has Dynamic Port	Indicates whether the service has dynamic ports or not.

Properties for vSAN

vRealize Operations Manager displays object properties for vSAN.

Properties for vSAN Disk Groups

vRealize Operations Manager displays the following property for vSAN disk groups:

- vSAN Disk Groups: Configuration|vSAN Configuration
- vSAN Disk Groups: Configuration | Number of Disks

Properties for vSAN Cluster

The vRealize Operations Manager displays the following properties for vSAN cluster.

Property Name	Description
Configuration vSAN Deduplication and Compression Enabled	Indicates whether deduplication and compression is enabled on the vSAN cluster.
Configuration vSAN Preferred fault domain	Indicates whether the preferred fault domain is not set for the witness host in a vSAN Stretched cluster.
Configuration vSAN Stretched Cluster	Indicates whether vSAN stretch cluster is enabled or not.
Configuration vSAN vSAN Configuration	Indicates whether the vSAN cluster is configured or not.
Configuration vSAN Encryption	Indicates whether the vSAN cluster is encrypted or not.
Configuration vSAN File Service	Indicates whether vSAN File Services is enabled or not.
Configuration vSAN File Service Domain:<domainName> DNS Servers	Indicates the IP addresses of DNS servers, which are used to resolve the host names within the DNS domain.
Configuration vSAN File Service Domain:<domainName> DNS Suffixes	Indicates the list of DNS suffixes which can be resolved by the DNS servers.
Configuration vSAN File Service Domain:<domainName> Gateway	Indicates the default gateway IP address for the file service access point.
Configuration vSAN File Service Domain:<domainName> Primary IP	Indicates the primary IP address for the file service.
Configuration vSAN File Service Domain:<domainName> Subnet Mask	Indicates the subnet mask for the vSAN cluster.
Summary Type	vSAN Cluster Type
Configuration vSAN File Service Domain:<domainName> IP Address :<ipaddress> FQDN	Indicates the Full Qualified Domain name (FQDN) to be used with IP address for the vSAN File Server instance.

Properties for vSAN Enabled Host

The vRealize Operations Manager displays the following property for vSAN enabled host.

- Configuration|vSAN Enabled
- Configuration|vSAN|Encryption

Properties for vSAN Cache Disk

vRealize Operations Manager displays the following properties for the vSAN cache disk.

Properties for vSAN include:

Component	Metrics
Configuration	<ul style="list-style-type: none"> ■ Configuration Properties Name ■ Configuration Properties Size ■ Configuration Properties Vendor ■ Configuration Properties Type ■ Configuration Properties Queue Depth ■ Configuration vSAN Encryption ■ Configuration Model
SCSI SMART Statistics	<ul style="list-style-type: none"> ■ 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

vRealize Operations Manager displays the following properties for the vSAN capacity disk.

Properties for vSAN include:

Component	Metrics
Configuration	<ul style="list-style-type: none"> ■ Configuration Properties Name ■ Configuration Properties Size ■ Configuration Properties Vendor ■ Configuration Properties Type ■ Configuration Properties Queue Depth ■ Configuration vSAN Encryption
SCSI SMART Statistics	<ul style="list-style-type: none"> ■ 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 File Server

The vRealize Operations Manager displays the following properties for vSAN file server.

- Configuration | vSAN | Primary
- Configuration | vSAN | FQDN

Properties for vSAN File Share

The vRealize Operations Manager displays the following properties for vSAN file share.

- Configuration |vSAN| Domain Name
- Configuration | vSAN| Hard Quota
- Configuration |vSAN| Soft Quota
- Configuration |vSAN | Label|<key>
- Configuration |vSAN | Access Point|<key>
- Configuration | vSAN | Permission:<permission> | Client IP Range
- Configuration | vSAN | Permission:<permission> | Root Squash

Properties for vRealize Automation 8.x

vRealize Operations Manager displays properties for vRealize Automation 8.x objects.

Some of the useful properties for project objects deployed through vRealize Automation 8.x are as follows:

- Project|CustomProperties: Custom properties defined for the project.

- Project|OrganizationID: Organization ID of the project.
- Project|userEmail: Email address of the user for the project.

One of the useful properties for the deployment object is:

- Deployment|User: User associated with the deployment.

One of the useful properties for the cloud zone object is:

- CloudAutomation|ResourceTags: Resource tags associated with the cloud zone.

One of the useful properties for the blueprint object is:

- Blueprint|User: User associated with the blueprint.

One of the useful properties for the CASworkd object is:

- CASWorld|metering|MeteringPolicyId: Metering policy ID associated with the CAS World object.

One of the useful properties for the virtual machine object is:

- Cloud Automation|CustomProperties: Custom properties associated with the virtual machine.

One of the useful properties for Cloud Zone is:

- Cloud Automation|Resource Tags: Resources tags associated with the cloud automation.

Properties in the NSX-T Adapter

vRealize Operations Manager displays the following properties for the NSX-T adapter.

Table 8-224. Properties in the NSX-T Adapter

Resource	Properties common in NSX-T and NSX-T on VMware Cloud on AWS	Properties in NSX-T on-premise	Properties NSX-T on VMware Cloud on AWS
Management Cluster		<ul style="list-style-type: none"> ■ NSXT Product Version ■ Status Summary Cluster Status Management Cluster Status ■ Status Summary Cluster Status Controller Cluster Status ■ Status Summary vIDM Connection Status ■ Status Summary Compute Managers <ComputeManagerName> Status ■ Configuration Maximums <ul style="list-style-type: none"> ■ Compute Manager count ■ Prepared vC Cluster count 	
Firewall Section	Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned Configuration <ul style="list-style-type: none"> ■ Firewall Rule Count Size 	Configuration <ul style="list-style-type: none"> ■ Firewall Stateful 	Configuration <ul style="list-style-type: none"> ■ Type ■ Domain id ■ Precedence ■ Category

Table 8-224. Properties in the NSX-T Adapter (continued)

Resource	Properties common in NSX-T and NSX-T on VMware Cloud on AWS	Properties in NSX-T on-premise	Properties NSX-T on VMware Cloud on AWS
Transport Node		<ul style="list-style-type: none"> ■ Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned ■ Summary FQDN ■ Status Summary <ul style="list-style-type: none"> ■ Transport Node State ■ Transport Node Deployment State ■ LCA Connectivity Status ■ Management Plane Connectivity Status ■ Host Node Deployment Status ■ Management connection Status ■ Controller connection Status ■ Load Balancer Usage <ul style="list-style-type: none"> ■ Current Small LB services ■ Current Medium LB services ■ Current Large LB services ■ Current Extra Large LB services ■ Current LB Pools ■ Current LB Pool Members ■ Current LB Virtual Servers ■ Remaining Small LB services ■ Remaining Medium LB services ■ Remaining Large LB services ■ Remaining Extra Large LB services ■ Remaining LB Pool Members ■ Tunnels <Tunnel-Name> Status ■ File Systems <FileSystemMount> <ul style="list-style-type: none"> ■ Total ■ Type ■ File System ID 	

Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.

Table 8-224. Properties in the NSX-T Adapter (continued)

Resource	Properties common in NSX-T and NSX-T on VMware Cloud on AWS	Properties in NSX-T on-premise	Properties NSX-T on VMware Cloud on AWS
Load Balancer Service <hr/> Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.		<ul style="list-style-type: none"> ■ Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned ■ LB Service Operational Status 	
Load Balancer Virtual Server <hr/> Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.		<ul style="list-style-type: none"> ■ Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned ■ LB Virtual Operational State 	
Load Balancer Pool <hr/> Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.		<ul style="list-style-type: none"> ■ Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned ■ Status 	
Transport Zone <hr/> Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.		Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ Switch Mode ■ System Owned 	

Table 8-224. Properties in the NSX-T Adapter (continued)

Resource	Properties common in NSX-T and NSX-T on VMware Cloud on AWS		Properties in NSX-T on-premise	Properties NSX-T on VMware Cloud on AWS
Logical Router	■	<ul style="list-style-type: none"> Summary <ul style="list-style-type: none"> Create Time Create User Last Modified Time Last Modified User Protection Revision System Owned 	<ul style="list-style-type: none"> Configuration <ul style="list-style-type: none"> Failover Mode High Availability Mode Edge Cluster Id Router Type Services Enabled <ul style="list-style-type: none"> HA Status Per Transport Node <TransportNodeID> HA Status Firewall Enabled Load balancer Enabled DNS Enabled L2VPN Enabled IPSEC VPN Enabled 	
Router Service	1	Tier-0 Router Services → BGP Service <ul style="list-style-type: none"> Summary BGP Neighbor Count 	<ul style="list-style-type: none"> All logical routers → Static Routes → Summary Static Route Count All logical routers → NAT Rule → Summary NAT Rule Count 	
	2	Tier-1 Router Services → NAT Rules <ul style="list-style-type: none"> Summary NAT Rule Count 	<ul style="list-style-type: none"> Tier 0 → BGP Service → Summary <ul style="list-style-type: none"> ECMP Status Status 	
	3	Tier-1 Router Services → Static Routes <ul style="list-style-type: none"> Summary Static Route Count 	<ul style="list-style-type: none"> Tier 0 → BFD Service → Summary <ul style="list-style-type: none"> Status BFD Neighbor Count Tier 0 → Route Redistribution → Summary <ul style="list-style-type: none"> Status Redistribution Rule count Tier 1 → Route Advertisement → Summary <ul style="list-style-type: none"> Route Advertisement Count Status 	
Logical Switch	■	<ul style="list-style-type: none"> Summary <ul style="list-style-type: none"> Create Time Create User Last Modified Time Last Modified User Protection Revision System Owned 	<ul style="list-style-type: none"> Summary <ul style="list-style-type: none"> Logical Switch State Configuration <ul style="list-style-type: none"> Replication Mode Admin State VNI 	Configuration <ul style="list-style-type: none"> Type

Table 8-224. Properties in the NSX-T Adapter (continued)

Resource	Properties common in NSX-T and NSX-T on VMware Cloud on AWS	Properties in NSX-T on-premise	Properties NSX-T on VMware Cloud on AWS
Management Appliances		NSXT API Version	
Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.			
Manager Node		<ul style="list-style-type: none"> ■ NSXT Manager Node Version ■ Connectivity Status Management Plane Connectivity Status 	
Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.			
Group	Configuration Maximums Count <ul style="list-style-type: none"> ■ IP Address Count ■ Expressions Count ■ vm Count 	Configuration Maximums Count Tag Count	
Edge Cluster		Summary <ul style="list-style-type: none"> ■ Create Time ■ Create User ■ Last Modified Time ■ Last Modified User ■ Protection ■ Revision ■ System Owned ■ Edge Cluster Member Type 	
Note This object is specific to NSX-T on-premise and is not available in NSX-T on VMware Cloud on AWS.			

Placement Group Properties

The following properties are available for each Placement Group instance in your vRealize Operations Manager environment.

Table 8-225. Placement Group Properties

Service	Property
Placement Group	State
	Strategy

Properties for VeloCloud Gateway

vRealize Operations Manager displays properties of VeloCloud Gateway objects.

Some of the useful properties for VeloCloud Gateway are as follows:

- Summary | Core Count
- Summary | Gateway Activation Status
- Summary | Gateway Network Interface Errors
- Summary | Gateway Time Zone
- Summary | ICMP Status
- Summary | Is Eth0 DPDK Enabled
- Summary | Is Eth1 DPDK Enabled
- Summary | Registration Status
- Summary | VCO IP
- Summary | Version

Properties for VeloCloud Orchestrator

vRealize Operations Manager displays properties of VeloCloud Orchestrator objects.

Some of the useful properties for VeloCloud Orchestrator are as follows:

- General | DR SSH Tunnel Status
- General | Internet Connectivity
- General | IP Address
- General | NTP Time Zone