

Definitions for Metrics, Properties, and Alerts

17 MAY 2022

vRealize Operations 8.6

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

About vRealize Operations Reference for Metrics, Properties, and Alerts 8

1 Metric Definitions in vRealize Operations 9

Metrics for vCenter Server Components 10

vSphere Metrics 10

vCenter Server Metrics 18

Virtual Machine Metrics 24

Host System Metrics 49

Cluster Compute Resource Metrics 73

Resource Pool Metrics 88

Data Center Metrics 91

Custom Data Center Metrics 100

Storage Pod Metrics 105

VMware Distributed Virtual Switch Metrics 106

Distributed Virtual Port Group Metrics 108

Datastore Cluster Metrics 110

Datastore Metrics 110

Cluster Compute Metrics for Allocation Model 118

Virtual Machine Metrics for Allocation Model 118

Metrics for Namespace 119

Metrics for Tanzu Kubernetes cluster 121

Metrics for vSphere Pods 122

OS and Application Monitoring Metrics 138

Operating System Metrics 138

Application Service Metrics 142

Windows Service Metrics 181

Linux Process Metrics 181

Remote Check Metrics 182

VeloCloud Application Service Metrics 183

Service Discovery Metrics 187

Virtual Machine Metrics 187

Service Summary Metrics 188

Service Performance Metrics 188

Service Type Metrics 188

Calculated Metrics 189

Capacity Analytics Generated Metrics 189

Badge Metrics 198

System Metrics 199

Log Insight Generated Metrics	200
Self-Monitoring Metrics for vRealize Operations	201
Analytics Metrics	201
Collector Metrics	206
Controller Metrics	206
FSDB Metrics	207
Product UI Metrics	207
Admin UI Metrics	208
Suite API Metrics	209
Cluster and Slice Administration Metrics	210
Watchdog Metrics	211
Node Metrics	212
Cluster Metrics	217
Persistence Metrics	224
Remote Collector Metrics	227
vRealize Automation 8.x Metrics	232
Blueprint Metrics	232
Project Metrics	232
Deployment Metrics	233
Organization Metrics	233
vRealize Adapter 8.x Metrics	234
Cloud Automation Services World Metrics	234
Cloud Automation Services Entity Status Metrics	234
Cloud Zone Limits Metrics	235
Metrics for vSAN	235
Disk I/O and Disk Space Metrics for vSAN Disk Groups	236
Read Cache Metrics for vSAN Disk Groups	236
Write Buffer Metrics for vSAN Disk Groups	237
Congestion Metrics for vSAN Disk Groups	237
Cache De-stage Metrics for vSAN Disk Groups	238
Resync Traffic Metrics for vSAN Disk Groups	238
Metrics for vSAN Cluster	238
Metrics for vSAN Enabled Host	240
Metrics for vSAN Datastore	241
Metrics for vSAN Cache Disk	241
Metrics for vSAN Capacity Disk	243
Metrics for vSAN Fault Domain Resource Kind	245
Metrics for vSAN World	246
Metrics for vSAN File Server	247
Metrics for vSAN File Share	247
Capacity Model for vSAN Objects	247

Metrics for Microsoft Azure	248
Virtual Machine Metrics	248
Cosmos DB Metrics	249
SQL Server Metrics	250
SQL Database Metrics	251
MySQL Server Metrics	254
PostgreSQL Server Metrics	255
Network Interface Metrics	256
Load Balancer Metrics	257
Metrics for Management Pack for AWS	257
EC2 Metrics	258
EC2 Volume Metrics	259
EC2 Load Balancer Metrics	260
Network Load Balancer Metrics	261
Application Load Balancer Metrics	261
EC2 Auto Scale Group Metrics	262
EMR Job Flow Metrics	263
Entity Status Metrics	265
ElastiCache Cache Node Metrics	266
RDS DB Instance Metrics	269
Lambda Metrics	269
Redshift Cluster Metrics	270
Redshift Node Metrics	270
AWS Workspace Metrics	271
ECS Cluster Metrics	272
ECS Service Metrics	272
DynamoDB Metrics	272
S3 Bucket Metrics	273
VPC Nat Gateway Metrics	274
Dax Cluster Metrics	275
DAX Node Metrics	276
Direct Connect Metrics	277
Health Check Metrics	278
ElastiCache Cache Cluster Metrics	278
EFS Metrics	279
Elastic Beanstalk Environment Metrics	280
AWS Transit Gateway Metrics	281
EKS Cluster Metrics	281
Metrics in VMware Cloud on AWS	282
Metrics in NSX-T Adapter	289
Metrics in VMware vRealize Operations Management Pack for Google Cloud Platform	297

Sustainability Metrics 308

2 Property Definitions in vRealize Operations 312

Properties for vCenter Server Components 313

vCenter Server Properties 313

Virtual Machine Properties 314

Host System Properties 321

Cluster Compute Resource Properties 326

Resource Pool Properties 327

Data Center Properties 328

Storage Pod Properties 329

VMware Distributed Virtual Switch Properties 329

Distributed Virtual Port Group Properties 329

Datastore Properties 330

vSphere Pod Properties 333

Namespace Properties 342

Tanzu Kubernetes cluster Properties 343

Self-Monitoring Properties for vRealize Operations 345

Analytics Properties 345

Node Properties 345

Remote Collector Properties 346

OS and Application Monitoring Properties 346

Guest Information Properties 346

Service Discovery Properties 346

Service Discovery Adapter Instance Properties 347

Virtual Machine Properties 347

Services Properties 347

Properties for vSAN 348

Properties for vSAN Disk Groups 348

Properties for vSAN Cluster 348

Properties for vSAN Enabled Host 349

Properties for vSAN Cache Disk 349

Properties for vSAN Capacity Disk 350

Properties for vSAN File Server 351

Properties for vSAN File Share 351

Properties for vRealize Automation 8.x 351

Properties in the NSX-T Adapter 352

Placement Group Properties 356

Properties for VeloCloud Gateway 357

Properties for VeloCloud Orchestrator 357

Sustainability Properties 357

3	Alert Definitions in vRealize Operations Manager	359
	Cluster Compute Resource Alert Definitions	360
	Host System Alert Definitions	365
	vRealize Automation Alert Definitions	378
	vSAN Alert Definitions	380
	Alerts in the vSphere Web Client	392
	vSphere Distributed Port Group	393
	Virtual Machine Alert Definitions	393
	vSphere Distributed Switch Alert Definitions	400
	vCenter Server Alert Definitions	401
	Datastore Alert Definitions	403
	Data Center Alert Definitions	408
	Custom Data Center Alert Definitions	409
	vSphere Pod Alert Definitions	410
	VMware Cloud on AWS Alert Definitions	414
	Alerts in VMware vRealize Operations Management Pack for Google Cloud Platform	417

About vRealize Operations Reference for Metrics, Properties, and Alerts

The *vRealize Operations Reference for Metrics, Properties, and Alerts* provides information about the metric, properties, and alert definitions provided with vRealize Operations.

Intended Audience

This information is intended for anyone who wants to install and configure vRealize Operations by using a virtual appliance deployment. The information is written for experienced virtual machine administrators who are familiar with enterprise management applications and data center operations.

Note All unit conversions in vRealize Operations are based on 1024 factor.

Metric Definitions in vRealize Operations

1

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 to display results that help you to manage your environment.

vRealize Operations collects data from objects in your environment. Each piece of data collected is called a metric observation or value. vRealize Operations uses the VMware vCenter® adapter to collect raw metrics. vRealize Operations uses the vRealize Operations adapter to collect self-monitoring metrics. In addition to the metrics it collects, vRealize Operations 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.

This chapter includes the following topics:

- [Metrics for vCenter Server Components](#)
- [OS and Application Monitoring Metrics](#)
- [Service Discovery Metrics](#)
- [Calculated Metrics](#)
- [Self-Monitoring Metrics for vRealize Operations](#)
- [vRealize Automation 8.x Metrics](#)
- [Metrics for vSAN](#)
- [Metrics for Microsoft Azure](#)
- [Metrics for Management Pack for AWS](#)
- [Metrics in VMware Cloud on AWS](#)
- [Metrics in NSX-T Adapter](#)
- [Metrics in VMware vRealize Operations Management Pack for Google Cloud Platform](#)
- [Sustainability Metrics](#)

Metrics for vCenter Server Components

vRealize Operations 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 1-1. 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 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](#).

vSphere World Super Metrics for ROI Dashboard

vSphere world super metrics provide information about the new metrics added to the ROI dashboard.

Metric Name	Description
Cost Total Cost of Ownership	This metric shows the total cost cost of ownership with potential savings and optimizations. Key: cost total_aggregated_cost
Online Capacity Analytics Capacity Remaining Profiles	This metric shows the VMs remaining based on the average VM profile. Key: OnlineCapacityAnalytics capacityRemainingProfile
Cost Server Hardware(Owned) Cost	This metric shows the sum of server hardware depreciated cost having purchase type as Owned across all the vCenters. Key: cost total_serverHardware_owned_cost
Cost Server Hardware(Leased) Cost	This metric shows the sum of server hardware depreciated cost having purchase type as Leased across all the vCenters. Key: cost total_serverHardware_leased_cost
Cost Host OS License cost	This metric shows the sum of host OS license cost across all the vCenters. Key: cost total_hostOsl_cost
Cost Network Cost	This metric shows the sum of network cost across all the vCenters. Key: cost total_network_cost
Cost Maintenance Cost	This metric shows the sum of maintenance cost across all the vCenters. Key: cost total_maintenance_cost
Cost Server Labor Cost	This metric shows the sum of server labor cost across all the vCenters. Key: cost total_serverLabor_cost
Cost Facilities Cost	This metric shows the sum of facilities cost across all the vCenters. Key: cost total_facilities_cost
Cost Additional Cost	This metric shows the sum of additional cost across all the vCenters. Key: cost total_additional_cost
Cost VM Direct Cost	This metric shows sum of direct Cost (VI labor + OS Labor) across all vCenters. Key: cost total_vm_direct_cost
Cost Capacity Used Compute Cost	This metric displays the cost of the used compute capacity. Key: cost capacity_used compute
Cost Capacity Remaining Compute Cost	This metric displays the cost of the remaining compute capacity. Key: cost capacity_remaining compute
Cost Capacity Used Storage Cost	This metric displays the cost of the used storage capacity. Key: cost capacity_used storage
Cost Capacity Remaining Storage Cost	This metric displays the cost of the remaining storage capacity. Key: cost capacity_remaining storage
Cost Potential Savings Idle VMs	This metric displays the potential savings from Idle VMs. Key: cost potential_savings idle_vms

Metric Name	Description
Cost Potential Savings Powered Off VMs	This metric displays the potential savings from powered off VMs. Key: cost potential_savings poweredOff_vms
Cost Potential Savings VM Snapshots	This metric displays the potential savings from VM snapshots. Key: cost potential_savings vm_snapshots
Cost Potential Savings Orphaned Disks	This metric displays the potential savings from orphaned disks. Key: cost potential_savings orphaned_disks
Cost Potential Savings Oversized VMs	This metric displays the potential savings from oversized VMs. Key: cost potential_savings oversized_vms
Cost Potential Savings Cost Optimization Opportunities	This metric displays the potential savings from cost optimization opportunities. Key: cost potential_savings cost_optimization_opportunities
Cost Total Cost of Ownership	This metric shows the total cost cost of ownership with potential savings and optimizations. Key: cost potential_savings total_cost_of_ownership
Server Purchase Cost	This metric shows the server purchase cost. Key: cost server_purchase_cost
Accumulated Depreciation	This metric displays the sum of the accumulated depreciation (Depreciation is calculated from the purchase date till current date) of servers across all the vCenters. Key: cost accumulatedDepreciation
Remaining Depreciation	This metric displays the sum of the remaining depreciation (Remaining Depreciation is calculated from the current date till Depreciated year) of servers across all the vCenters. Key: cost accumulatedDepreciation
Number of Fully Depreciated Servers	This metric displays the number of fully depreciated servers across all the vCenters. Key: cost hardwareTotalCost
Reclaimed vCPUs from Idle VMs	This metric displays the number of reclaimable vCPUs from idle VMs. Key: reclaimable idle_vms cpu
Reclaimed Memory from Idle VMs	This metric displays the amount of reclaimable memory from the idle VMs. Key: reclaimable idle_vms mem
Reclaimed Disk Space from Idle VMs	This metric displays the amount of reclaimable disk space from the idle VMs. Key: reclaimable idle_vms diskspace
Reclaimed Disk Space from Powered Off VMs	This metric displays the amount of reclaimable disk space from the powered off VMs. Key: reclaimable poweredOff_vms diskspace
Reclaimed Disk Space from VM Snapshots	This metric displays the amount of reclaimable disk space from the VM Snapshots. Key: reclaimable vm_snapshots diskspace
Reclaimed Disk Space from Orphaned Disks	This metric displays the amount of reclaimable disk space from the orphaned disks. Key: reclaimable orphaned_disk diskspace

Metric Name	Description
Rightsize - vCPUs to Remove from Oversized VMs	This metric displays the number of vCPUs to remove from the oversized VMs. Key: summary oversized vcpus
Rightsize - Memory to Remove from Oversized VMs	This metric displays the amount of memory to be removed from the oversized VMs. Key: summary oversized memory
Rightsize - vCPUs to Add from Undersized VMs	This metric displays the number of vCPUs to be added from the undersized VMs. Key: summary undersized vcpus
Rightsize - Memory to Add from Undersized VMs	This metric displays the amount of memory to be added from the undersized VMs. Key: summary undersized memory
Total Storage Cost	This metric displays the sum of storage cost across all vCenters. Key: cost totalCost
Total Potential Savings	This metric displays the sum of all the potential savings (Idle VMs + Powered off Vms + Snapshot + Orphaned Disks + Oversized VMs). Key: reclaimable cost
New vSphere Metrics Added for ROI Dashboard	
Potential Savings from Oversized VMs	This metric displays the sum of all the potentials savings gained from oversized VMs across vcenters. Key: cost reclaimableCost
Reclaimable Host Cost	This metric displays the reclaimable host cost based on the recommended size. Key: cost potential_savings total_reclaimable_host_cost
Cost Potential Increase Undersized VMs Cost	This metric displays the rightsizing value for the undersized VMs. Key: cost potential_increase undersized_vms
Cost Realized Savings Total Realized Savings	This metric displays the total realize savings for VMs across all vCenters. Key: cost realized_savings total_realized_savings
Cost Realized Savings Idle Savings	This metric displays the total realized savings for idle VMs across all vCenters. Key: cost realized_savings realized_idle_savings
Cost Realized Savings Powered Off Savings	This metric displays the total realized savings for powered off VMs across all vCenters. Key: cost realized_savings realized_poweredOff_savings
Cost Realized Savings Snapshot Space Savings	This metric displays the total realized savings for snapshot space across all vCenters. Key: cost realized_savings realized_snapshotSpace_savings
Cost Realized Savings Oversized Savings	This metric displays the oversized savings across all vCenters. Key: cost realized_savings realized_oversized_savings
Cost Realized Savings Orphaned Disk Space Savings	This metric displays the amount of disk space saved by orphaned disks across all vCenters. Key: cost realized_savings realized_orphanedDiskSpace_savings

Metric Name	Description
Cost Realized Savings Reclaimable Host Savings	This metric displays the amount of reclaimable host savings across all vCenters. Key: cost realized_savings realized_reclaimableHost_savings
Compute Realized vCPUs from Oversized VMs	This metric displays the number of vCPUs realized across all vCenters. Key: compute_realized realized_oversized_vcpus
Compute Realized Memory from Oversized VMs	This metric displays the amount of memory realized from oversized VMs across all vCenters. Key: compute_realized realized_oversized_mem
Realized Potential Memory Consumed from Oversized VMs	This metric displays the potential memory consumed from oversized VMs across all vCenters. Key: realized realizedPotentialMemConsumed
Total Number Of Reclaimable Hosts	This metric displays the total number of reclaimable hosts across all vCenters. Key: metric=cost reclaimableHostCost
Compute Realized vCPUs from Idle VMs	This metric displays the realized vCPUs from idle VMs across all vCenters. Key: compute_realized realized_idle_vcpus
Compute Realized Memory from Idle VMs	This metric displays the amount of memory realized from idle VMs across all vCenters. Key: compute_realized realized_idle_mem
Disk Space Realized Idle VMs	This metric displays the amount of disk space realized from idle VMs across all vCenters. Key: storage_realized realized_idle_diskSpace
Disk Space Realized Powered Off VMs	This metric displays the amount of disk space realized from powered off VMs across all vCenters. Key: storage_realized realized_poweredOff_diskSpace
Disk Space Realized VM Snapshots	This metric displays the amount of disk space realized from VM snapshots across all vCenters. Key: storage_realized realized_snapshotSpace
Disk Space Realized Orphaned Disks	This metric displays the amount of disk space realized from orphaned disks across all vCenters. Key: storage_realized realized_orphaned_diskSpace

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
summary Number of Reclaimable Hosts	Displays the number of reclaimable hosts. Key: summary total_number_reclaimable_hosts

vCenter Server Metrics

vRealize Operations 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: 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

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: diskspace total_usage
Total disk space (KB)	Total disk space on all datastores visible to this object. Key: diskspace total_capacity
Total provisioned disk space (KB)	Total provisioned disk space on all datastores visible to this object. Key: diskspace total_provisioned
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

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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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 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.

Application Level Cost Roll up Metrics

The application-level cost roll up in vRealize Operations includes few additional metrics at the application level. vRealize Operations has introduced Business Application as a new object, the business application object can have Tiers and Applications as its children. The cost roll up option allows you to aggregate all VM costs associated with the application and tiers and publish them at VM level.

Table 1-2. Application Level Cost Roll up Metrics

Metric Name	Description
Monthly Effected Projected Total Cost	This metric displays the effective virtual machine cost projected for the full month.
Effective MTD Cost	This metric displays the month to date effective application cost for the selected virtual machine.
Effective Daily Cost	This metric displays the effective daily cost of the application associated with the virtual machine.

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: cpu host_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 Workload %	This metric indicates the CPU workload % for the VM, the maximum threshold for this is 80% and the minimum threshold is 20%. If your Maximum line is constantly ~100% flat, you may have a runaway process. If this chart is below or less than 20% all the time for the entire month, then all the large VMs are oversized. This number must hover around 40%, indicating the sizing done was accurate.
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

Metric Name	Description
CPU Peak vCPU Ready	The highest CPU Ready among the virtual CPUs. Key: cpu peak_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: cpu peak_vcpu_usage
CPU 20-second Peak CPU System (%)	The highest system CPU, 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</p> <p>rescpu actav5_latest</p> <p>rescpu actav15_latest</p> <p>rescpu actpk1_latest</p> <p>rescpu actpk5_latest</p> <p>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</p> <p>rescpu runav5_latest</p> <p>rescpu runav15_latest</p> <p>rescpu runpk1_latest</p> <p>rescpu runpk5_latest</p> <p>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</p> <p>rescpu maxLimited5_latest</p> <p>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 ISwapInRate_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 Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem ISwapUsed_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
Virtual Disk Total Throughput	Amount of data read from/written to storage in a second. This is averaged over the reporting period. 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

Metric Name	Description
Storage Total Usage (KBps)	Total throughput rate in kilobytes per second. Key: storage usage_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

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 . 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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

Metric Name	Description
CPU 50% of Recommended number of vCPUs to Remove	This metric is superseded by the capacity engine. cpulnumberToRemove50Pct
CPU Capacity entitlement (mhz)	cpulcapacity_entitlement
CPU Co-stop (msec)	Use the Co-Stop (%) metric instead of this metric. cpulcostop_summation
CPU Demand Over Capacity (mhz)	cpuldemandOverCapacity
CPU Demand Over Limit (mhz)	Use Contention (%) metric instead of this metric. cpuldemandOverLimit
CPU Dynamic entitlement	cpuldynamic_entitlement
CPU Estimated entitlement	cpulestimated_entitlement
CPU Idle (%)	cpulidlePct
CPU Idle (msec)	cpulidle_summation
CPU IO Wait (msec)	cpulawait
CPU Normalized Co-stop (%)	Use the Co-Stop (%) metric instead of this metric. cpulperCpuCoStopPct
CPU Provisioned vCPU(s) (Cores)	cpulcorecount_provisioned
CPU Ready (msec)	Choose the Use Ready (%) metric instead of this metric. cpulready_summation
CPU Recommended Size Reduction (%)	cpulsizePctReduction
CPU Swap Wait (msec)	cpulswapwait_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 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](#).

Host System Metrics for ROI Dashboard

Host system metrics provide information about cost saving across vCenters

Metric Name	Description
Cost Monthly Additional Total Cost	This metric shows the total sum of additional cost across all the vCenters for an entire month. Key: cost additionalTotalCost

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
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

Metric Name	Description
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: cpulcapacity_contentionPct</p>
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: cpuldemandPct</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: cpuldemandmhz</p>
CPU IO Wait (ms)	<p>IO wait time in milliseconds.</p> <p>Key: cpulawait</p>
CPU Number of CPU Sockets	<p>Number of CPU sockets.</p> <p>Key: cpulnumpackages</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>

Metric Name	Description
CPU Demand	CPU demand. Key: cpuldemand_average
CPU Usage (MHz)	CPU use in megahertz. Key: cpulusagemhz_average
CPU Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool. Key: cpulreservedCapacity_average
CPU Total Capacity (MHz)	Total CPU capacity in megahertz. Amount of CPU resources configured on the ESXi hosts. Key: cpulcapacity_provisioned
CPU Overhead (KB)	Amount of CPU overhead. Key: cpuoverhead_average
CPU Demand without overhead	Value of demand excluding any overhead. Key: cpuldemand_without_overhead
CPU Core Utilization (%)	Percent core utilization. Key: cpulcoreUtilization_average
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: cpu ready_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

Metric Name	Description
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>)	Average active time for the CPU over the past minute, past five minutes, and at one-minute, five-minute, and 15-minute peak active times. Key: rescpulactav1_latest rescpulactav5_latest rescpulactav15_latest rescpulactpk1_latest rescpulactpk5_latest rescpulactpk15_latest
Rescpu CPU Running (%) (<i>interval</i>)	Average run time for the CPU over the past minute, past five minutes, past 15 minutes, and at one-minute, five-minute, and 15-minute peak times. Key: rescpu runav1_latest rescpu runav5_latest rescpu runav15_latest rescpu runpk1_latest rescpu runpk5_latest rescpu runpk15_latest
Rescpu CPU Throttled (%) (<i>interval</i>)	Scheduling limit over the past minute, past five minutes, and past 15 minutes. Key: rescpu maxLimited1_latest rescpu maxLimited5_latest rescpu maxLimited15_latest
Rescpu Group CPU Sample Count	Group CPU sample count. Key: rescpu sampleCount_latest
Rescpu Group CPU Sample Period (ms)	Group CPU sample period in milliseconds. Key: rescpu samplePeriod_latest

Datastore Metrics for Host Systems

Datastore metrics provide information about datastore use.

Metric Name	Description
Datastore Outstanding IO requests	OIO for datastore. Key: datastore demand_oio
Datastore Commands Averaged	Average number of commands issued per second during the collection interval. Key: datastore commandsAveraged_average
Datastore Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore oio
Datastore 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

Metric Name	Description
Datastore Max Queue Depth	Max Queue Depth. Key: datastore datastoreMaxQueueDepth_latest
Datastore Highest Latency	Highest Latency. Key: datastore maxTotalLatency_latest
Datastore Total Latency Max	Total Latency Max (ms). Key: datastore totalLatency_max
Datastore Read Latency	Read Latency. Key: datastore datastoreNormalReadLatency_latest
Datastore Write Latency	Write Latency. Key: datastore datastoreNormalWriteLatency_latest
Datastore Data Read	Data Read. Key: datastore datastoreReadBytes_latest
Datastore Data Read Rate	Data Rate. Key: datastore datastoreReadIops_latest
Datastore Read Load	Storage DRS metric read load. Key: datastore datastoreReadLoadMetric_latest
Datastore Outstanding Read Requests	Outstanding Read Requests. Key: datastore datastoreReadOIO_latest
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
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

Metric Name	Description
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
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

Metric Name	Description
Mem Minimum Free Memory (KB)	Minimum free memory. Key: mem host_minfree
Mem Reserved Capacity (%)	Percent reserved capacity. Key: mem reservedCapacityPct
Mem Usable Memory (KB)	Usable memory in kilobytes. Key: mem host_usable
Mem Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem host_usagePct
Mem ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem host_systemUsage
Mem Guest Active (KB)	Amount of memory that is actively used. Key: mem active_average
Mem Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory. Key: mem consumed_average
Mem Granted (KB)	Amount of memory available for use. Key: mem granted_average
Mem Heap (KB)	Amount of memory allocated for heap. Key: mem heap_average
Mem Heap Free (KB)	Amount of free space in the heap. Key: mem heapfree_average
Mem VM Overhead (KB)	Memory overhead reported by host. Key: mem overhead_average
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

Metric Name	Description
Mem Unreserved (KB)	Amount of unreserved memory in kilobytes. Key: mem unreserved_average
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.</p> <p>The presence of ballooning indicates the ESXi has been under memory pressure. The ESXi activates ballooning when consumed memory reaches a certain threshold.</p> <p>Look for increasing size of ballooning. This indicates that there has been a shortage of memory more than once.</p> <p>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)	Amount of memory that is all zero. Key: mem zero_average
Mem State (0-3)	Overall state of the memory. The value is an integer between 0 (high) and 3 (low). Key: mem state_latest
Mem Usage (KB)	Host memory use in kilobytes. Key: mem host_usage
Mem Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem usage_average
Mem Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval in kilobyte per second. Key: mem swapiRate_average
Mem Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory to disk during the current interval in kilobytes per second. Key: mem swapoutRate_average
Mem Active Write (KB)	Average active writes in kilobytes. Key: mem activewrite_average
Mem Compressed (KB)	Average memory compression in kilobytes. Key: mem compressed_average
Mem Compression Rate (KBps)	Average compression rate in kilobytes per second. Key: mem compressionRate_average
Mem Decompression Rate (KBps)	Decompression rate in kilobytes per second. Key: mem decompressionRate_average

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: ysl 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	Average number of write commands issued per second by the storage adapter during the collection interval. Key: storageAdapter numberWriteAveraged_average
Storage Adapter Read Throughput (KBps)	Rate of reading data by the storage adapter. Key: storageAdapter read_average
Storage Adapter Read Latency (ms)	This metric shows the average amount of time for a read operation by the storage adapter. Use this metric to monitor the storage adapter read operation performance. A high value means that the ESXi is performing a slow storage read operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter totalReadLatency_average
Storage Adapter Write Latency (ms)	This metric shows the average amount of time for a write operation by the storage adapter. Use this metric to monitor the storage adapter write performance operation. A high value means that the ESXi is performing a slow storage write operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter totalWriteLatency_average
Storage Adapter Write Throughput (KBps)	Rate of writing data by the storage adapter. Key: storageAdapter write_average
Storage Adapter Demand	Demand. Key: storageAdapter demand
Storage Adapter Highest Latency	Highest Latency. Key: storageAdapter maxTotalLatency_latest
Storage Adapter Outstanding Requests	Outstanding Requests. Key: storageAdapter outstandingIOs_average
Storage Adapter Queue Depth	Queue Depth. Key: storageAdapter queueDepth_average
Storage Adapter Queue Latency (ms)	The average time spent in the ESX Server VM Kernel queue per command. Key: storageAdapter queueLatency_average
Storage Adapter Queued	Queued. Key: storageAdapter queued_average

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 hostOs TotalCost
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 . 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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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 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](#)

Cluster Metrics for ROI Dashboard

Cluster metrics provide information about the metrics in ROI dashboard.

Metric Name	Description
Total Number Of Reclaimable Hosts	<p>This metric displays the total number of reclaimable hosts across all vCenters.</p> <p>Key: metric=cost reclaimableHostCost</p>
Total Reclaimable Host Cost	<p>This metric displays the reclaimable host cost based on the recommended size.</p> <p>Key: cost reclaimableHostCost</p>

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>

Metric Name	Description
Configuration DRS Configuration Tolerance Imbalance Threshold	Displays the tolerance imbalance threshold for DRS configuration. Key: configuration DRSconfiguration ToleranceimbalanceThreshold
Configuration DRS Configuration Default DRS behavior	Displays the default DRS configuration behavior. Key: configuration DRSconfiguration DefaultDRSbehaviour
Configuration DRS Configuration Idle Consumed Memory	Displays the idle memory consumed by DRS configuration. Key: configuration DRSconfiguration IdleConsumedMemory
Configuration DRS Configuration DRS vMotion Rate	Displays the vMotion rate for the DRS configuration. Key: configuration DRSconfiguration DRSvMotion Rate
Configuration DPM Configuration Default DPM behavior	Displays the default behavior for the DPM configuration. Key: configuration DPMconfiguration DefaultDPMbehaviour
Configuration DPM Configuration DPM Enabled	Displays whehter 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

Metric Name	Description
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)	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. Key: cpu alloc usableCapacity
CPU Capacity Usage	This metric shows the percentage of the capacity used. Key: cpu capacity_usagepct_average

Metric Name	Description
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: <code>cpu capacity_contentionPct</code></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: <code>cpu demand usableCapacity</code></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: <code>cpu demandPct</code></p>
CPU Demand (MHz)	<p>Sum of CPU utilization of all virtual machines on this cluster, including limits and VM overhead.</p> <p>Key: <code>cpu demandmhz</code></p>
CPU Number of CPU Sockets	<p>Number of CPU sockets.</p> <p>Key: <code>cpu numpackages</code></p>
CPU Overall CPU Contention	<p>Overall CPU contention in milliseconds.</p> <p>Key: <code>cpu capacity_contention</code></p>
CPU Host Provisioned Capacity	<p>Provisioned CPU capacity in megahertz.</p> <p>Key: <code>cpu capacity_provisioned</code></p>
CPU Provisioned CPUs	<p>Number of Physical CPUs (Cores).</p> <p>Key: <code>cpu corecount_provisioned</code></p>

Metric Name	Description
CPU Usage (MHz)	Average CPU use in megahertz. Key: cpu usagemhz_average
CPU Demand	CPU Demand. Key: cpu demand_average
CPU Overhead	Amount of CPU overhead. Key: cpu overhead_average
CPU Demand without overhead	Value of demand excluding any overhead. Key: cpu demand_without_overhead
CPU Provisioned Capacity	Provisioned Capacity (MHz). Key: cpu vm_capacity_provisioned
CPU Number of hosts stressed	Number of hosts stressed. Key: cpu num_hosts_stressed
CPU Stress Balance Factor	Stress Balance Factor. Key: cpu stress_balance_factor
CPU Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: cpu min_host_capacity_remaining
CPU Workload Balance Factor	Workload Balance Factor. Key: cpu workload_balance_factor
CPU Highest Provider Workload	Highest Provider Workload. Key: cpu max_host_workload
CPU Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: cpu host_workload_disparity
CPU Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: cpu host_stress_disparity
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_average
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
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

Metric Name	Description
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
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

Metric Name	Description
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
Mem Usage (%)	This metric shows the portion of the total memory in all hosts in the cluster that is being used. This metric is the sum of memory consumed across all hosts in the cluster divided by the sum of physical memory across all hosts in the cluster. $\frac{\sum \text{memory consumed on all hosts}}{\sum \text{physical memory on all hosts}} - \times 100\%$
Mem Usage (KB)	Memory currently in use as a percentage of total available memory. Key: mem usage_average
Mem VM kernel Usage (KB)	Amount of memory that the VM kernel uses. Key: mem sysUsage_average
Mem Zero (KB)	Amount of memory that is all 0. Key: mem zero_average
Mem Number of Hosts Stressed	Number of hosts stressed. Key: mem num_hosts_stressed

Metric Name	Description
Mem Stress Balance Factor	Stress balance factor. Key: mem stress_balance_factor
Mem Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: mem min_host_capacity_remaining
Mem Workload Balance Factor	Workload balance factor. Key: mem workload_balance_factor
Mem Highest Provider Workload	Highest provider workload. Key: mem max_host_workload
Mem Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: mem host_workload_disparity
Mem Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: mem host_stress_disparity
Mem 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
Mem Total Capacity (KB)	Total physical memory configured on descendant ESXi hosts. Key: mem host_provisioned
Mem 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 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

Metric Name	Description
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
Network Error Packets	Displays the total number of error packets (transmitted and received) from all the ESXi in the cluster within a time interval of 20 seconds. Key: Network Error Packets

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
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

Metric Name	Description
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

Metric Name	Description
Summary Number of vMotions	<p>This metric shows the number of vMotions that occurred during the last collection cycle.</p> <p>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.</p> <p>Key: summary number_vmotion</p>
Summary Number of Hosts	<p>Total number of hosts.</p> <p>Key: summary total_number_hosts</p>
Summary Total Number of VMs	<p>Total number of virtual machines.</p> <p>Note This shows the total number of VMs excluding VM templates under the datastore.</p> <p>Key: summary total_number_vms</p>
Summary Total Number of Datastores	<p>Total number of datastores.</p> <p>Key: summary total_number_datastores</p>
Summary Number of VCPUs on Powered On VMs	<p>Number of virtual CPUs on powered-on virtual machines.</p> <p>Key: summary number_running_vcpus</p>
Summary Average Running VM Count per Running Host	<p>Average number of running virtual machines per running host.</p> <p>Key: summary avg_vm_density</p>
Summary Cluster Availability (%)	<p>Percentage of hosts powered-on in the cluster.</p> <p>Key: summary cluster_availability</p>
Summary Datastore	<p>Displays the status of the datastore.</p> <p>Key: summary datastore</p>
Summary Type	<p>Displays the datastore type.</p> <p>Key: summary type</p>
Summary Is Local	<p>Displays whether the datastore is local or not.</p> <p>Key: summary islocal</p>
Summary Number of VM Templates	<p>Number of VM templates.</p> <p>Key: summary number_vm_templates</p>
Summary Number of Pods	<p>Number of pods.</p> <p>Note This is published if the cluster is Workload Management enabled or there are pods under the cluster.</p> <p>Key: summary total_number_pods</p>
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>

Metric Name	Description
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 kuberntesStatus</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 diskspce</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>
Powered Off VMs Potential Savings	<p>Potential saving after reclamation of resources of Powered Off VMs within the cluster.</p> <p>Key: reclaimable poweredOff_vms cost</p>
VM Snapshots Disk Space (GB)	<p>Reclaimable disk space of VM Snapshots within the cluster.</p> <p>Key: reclaimable vm_snapshots diskspace</p>
VM Snapshots Potential Savings	<p>Potential saving after reclamation of VM Snapshots within the cluster.</p> <p>Key: reclaimable vm_snapshots cost</p>

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
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

Metric Name	Description
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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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

Metric Name	Key
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 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

Metric Name	Description
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
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 swapout_average

Metric Name	Description
Swap Used	Amount of memory used for swap space in kilobytes. Key: mem swapused_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 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](#)

Data Center Metrics for ROI Dashboard

Data center metrics provide information about data center savings across vCenters.

Metric Name	Description
Realized Cost Savings	
Realized Savings Idle Cost	This metric displays the total realize savings for VMs across all vCenters. Key: cost realized_savings realizedIdleCost
Realized Savings Powered Off Cost (AOA)	This metric displays the total realized savings for powered off VMs across all vCenters. Key: cost realized_savings realizedPoweredOffCost

Metric Name	Description
Realized Savings Snapshot Space Cost (AOA)	This metric displays the snapshots space saved across all vCenters. Key: cost realized_savings realizedSnapshotSpaceCost
Realized Savings Oversized Cost (AOA)	This metric displays the oversized savings across all vCenters. Key: cost realized_savings realizedOversizedCost
Realized Savings Orphaned Disk Space Cost (AOA)	This metric displays the amount of disk space saved by orphaned disks across all vCenters. Key: cost realized_savings realizedOrphanedDiskSpaceCost
Realized Savings Reclaimable Host Cost (AOA)	This metric displays the amount of reclaimable host savings across all vCenters. Key: cost realized_savings realizedReclaimableHostCost
Realized vCPUs from Oversized VMs	This metric displays the number of vCPUs realized across all vCenters. Key: realized realizedVCpus
Compute Realized Memory from Oversized VMs	This metric displays the amount of memory realized from oversized VMs across all vCenters. Key: compute_realized realizedOversizedMem
Realized Potential Memory Consumed from Oversized VMs	This metric displays the potential memory consumed from oversized VMs across all vCenters. Key: realized realizedPotentialMemConsumed
Compute Realized vCPUs from Oversized VMs	This metric displays the realized vCPUs from oversized VMs across all vCenters. Key: compute_realized realizedOversizedVCpus
Compute Realized vCPUs from Idle VMs	This metric displays the realized vCPUs from idle VMs across all vCenters. Key: compute_realized realizedIdleVCpus
Compute Realized Memory from Idle VMs	This metric displays the amount of memory realized from idle VMs across all vCenters. Key: compute_realized realizedIdleMem
Disk Space Realized Idle VMs	This metric displays the amount of disk space realized from idle VMs across all vCenters. Key: storage_realized realizedIdleDiskSpace
Disk Space Realized PoweredOff VMs	This metric displays the amount of disk space realized from powered off VMs across all vCenters. Key: storage_realized realizedPoweredOffDiskSpace
Disk Space Realized VM Snapshots	This metric displays the amount of disk space realized from VM snapshots across all vCenters. Key: storage_realized realizedSnapshotSpace

Metric Name	Description
Disk Space Realized Orphaned Disks	This metric displays the amount of disk space realized from orphaned disks across all vCenters. Key: storage_realized realizedIdleDiskSpace
Realized Savings Total Realized Cost	This metric displays the total realized cost across all vCenters. Key: cost realized_savings realizedTotalCost

CPU Usage Metrics for Data Centers

CPU usage metrics provide information about CPU use.

Metric Name	Description
Capacity Usage (%)	Percent capacity used. Key: cpulcapacity_usagepct_average
CPU Contention (%)	CPU capacity contention. Key: cpulcapacity_contentionPct
Demand (%)	CPU demand percentage. Key: cpu demandPct
Demand	Demand in megahertz. Key: cpuldemandmhz
Demand (MHz)	CPU utilization level based on descendant virtual machines utilization. This includes reservations, limits, and overhead to run the virtual machines. Key: cpu demandmhz
Overhead (KB)	Amount of CPU overhead. Key: cpu overhead_average
Demand without overhead	Value of demand excluding any overhead. Key: cpu demand_without_overhead
Total Wait	CPU time spent on idle state. Key: cpu wait
Number of CPU Sockets	Number of CPU sockets. Key: cpulnumpackages
Overall CPU Contention (ms)	Overall CPU contention in milliseconds. Key: cpulcapacity_contention
Host Provisioned Capacity (MHz)	Host provisioned capacity in megahertz. Key: cpu capacity_provisioned
Provisioned vCPU(s)	Provisioned vCPU(s). Key: cpulcorecount_provisioned
Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool. Key: cpulreservedCapacity_average

Metric Name	Description
Usage	Average CPU usage in megahertz. Key: cpu usagemhz_average
IO Wait	IO wait time in milliseconds. Key: cpu iowait
Provisioned Capacity	Provisioned Capacity. Key: cpu vm_capacity_provisioned
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	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: 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: 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

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

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

Metric Name	Description
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 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
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

Metric Name	Description
Utilization (GB)	Storage space used on the connected vSphere Datastores. Key: diskspaceltotal_usage
Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: diskspaceltotal_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
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

Metric Name	Description
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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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 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: cpulcapacity_provisioned
Provisioned vCPU(s)	Provisioned vCPU(s). Key: cpulcorecount_provisioned
Demand without overhead	Value of demand excluding any overhead. Key: cpuldemand_without_overhead
Number of hosts stressed	Number of hosts stressed. Key: cpulnum_hosts_stressed
Stress Balance Factor	Stress balance factor. Key: cpulstress_balance_factor
Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: cpulmin_host_capacity_remaining
Workload Balance Factor	Workload balance factor. Key: cpulworkload_balance_factor
Highest Provider Workload	Highest provider workload. Key: cpulmax_host_workload
Host workload Max-Min Disparity	Host workload max-min disparity. Key: cpulhost_workload_disparity
Host stress Max-Min Disparity	Difference of max and min host stress in the container. Key: cpulhost_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: cpuldemandmhz
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

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
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

Metric Name	Description
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 monitors multiple vCenters which use shared data stores.
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: diskspaceltotal_usage
Total Capacity (GB)	Total storage space available on the connected vSphere datastores. Key: diskspaceltotal_capacity

Disabled Metrics

The following metrics are disabled in this version of vRealize Operations . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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
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 observer number of outstanding IO operations. Key: datastore maxObserved_OIO

Storage Pod Metrics

vRealize Operations 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 1-3. 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

Table 1-3. Datastore Metrics for Storage Pods (continued)

Metric Name	Description
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
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 1-4. 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 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 1-5. 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
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 1-6. 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 1-7. 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 1-8. 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 1-8. 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 1-9. 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 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 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 . 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 . This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect 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 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 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 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 1-10. Metrics for Namespace

Metric Key	Localized Name	Description
cpulusagemhz_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).
cpuleffective_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.

Table 1-10. Metrics for Namespace (continued)

Metric Key	Localized Name	Description
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.
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 1-10. 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 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 1-11. 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).
cpu effective_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 1-11. 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 swpoutRate_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 swpout_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 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 swapiRate_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 swapi_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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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 1-12. 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

OS and Application Monitoring Metrics

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

Operating System Metrics

Metrics are collected for Linux and Windows operating systems.

Linux Platforms

The following metrics are collected for Linux operating systems:

Table 1-13. 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

Table 1-13. Metrics for Linux (continued)

Metric	Metric Category	KPI
<Instance name> Usage System (%)	CPU	True
<Instance name> Usage User (%)	CPU	True
CPU Load1 (%)	CPU Load	False
CPU Load15 (%)	CPU Load	False
CPU Load5 (%)	CPU Load	False
<Instance name> IO Time	Disk IO	False
<Instance name> Read Time	Disk IO	False
<Instance name> Reads	Disk IO	False
<Instance name> Write Time	Disk IO	False
<Instance name> Writes	Disk IO	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

Table 1-13. Metrics for Linux (continued)

Metric	Metric Category	KPI
Total	Swap	True
Used	Swap	True
Used Percent	Swap	True

Windows Platforms

The following metrics are collected for Windows operating systems:

Table 1-14. 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
DPC Time (%)	CPU	False
Usage Guest (%)	CPU	False
Usage System (%)	CPU	False
Usage User (%)	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 1-14. 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
Total (bytes)	Memory	False
Used (bytes)	Memory	False
Used Percent(%)	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

Table 1-14. Metrics for Windows (continued)

Metric	Metric Category	KPI
Threads	System	False
Used Percent (%)	Swap	False
Total (bytes)	Swap	False

Application Service Metrics

Metrics are collected for 23 application services.

Active Directory Metrics

Metrics are collected for the Active Directory application service.

Table 1-15. 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

Table 1-15. Active Directory Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-15. Active Directory Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	Active Directory	False

ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

Table 1-16. ActiveMQ Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Active MQ	False
Buffer Pool<InstanceName> Memory Used	Active MQ	False
Buffer Pool<InstanceName> Total Capacity	Active MQ	False
Class Loading Loaded Class Count	Active MQ	False
Class Loading Unloaded Class Count	Active MQ	False
Class Loading Total Loaded Class Count	Active MQ	False
File Descriptor Usage Max File Descriptor Count	Active MQ	False
File Descriptor Usage Open File Descriptor Count	Active MQ	False

Table 1-16. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	Active MQ	False
Threading Thread Count	Active MQ	False
Uptime	Active MQ	False
UTILIZATION Process CpuLoad	Active MQ	False
UTILIZATION Memory Limit	ActiveMQ Broker	True
UTILIZATION Memory Percent Usage (%)	ActiveMQ Broker	True
UTILIZATION Store Limit	ActiveMQ Broker	False
UTILIZATION Store Percent Usage (%)	ActiveMQ Broker	False

Table 1-16. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
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
UTILIZATION Process CpuLoad	ActiveMQ OS	False
UTILIZATION System Cpu Load	ActiveMQ OS	False

Table 1-16. ActiveMQ Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Consumer Count	ActiveMQ Topic	True
UTILIZATION Dequeue Count	ActiveMQ Topic	True
UTILIZATION Enqueue Count	ActiveMQ Topic	True
UTILIZATION Queue Size	ActiveMQ Topic	True
UTILIZATION Producer Count	ActiveMQ Topic	False

Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

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

Table 1-17. 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
UTILIZATION SCBoard Reading	Apache HTTPD	False
UTILIZATION SCBoard Sending	Apache HTTPD	False

Table 1-17. Apache HTTPD Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	Apache HTTPD	False

Apache HTTPD

Metrics are collected for the Apache HTTPD application service.

Table 1-18. Apache HTTPD

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

Table 1-18. Apache HTTPD (continued)

Metric Name	Category	KPI
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
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

Table 1-18. Apache HTTPD (continued)

Metric Name	Category	KPI
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
Application Availability	Tomcat Server	False
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
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

Microsoft IIS Metrics

Metrics are collected for the Microsoft IIS application service.

Table 1-19. Microsoft 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

Table 1-19. Microsoft IIS Metrics (continued)

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

Table 1-19. Microsoft IIS Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	Microsoft IIS	False

Java Application Metrics

Metrics are collected for the Java application service.

Table 1-20. 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

Table 1-20. Java Application Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-20. Java Application Metrics (continued)

Metric Name	Category	KPI
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 Server Metrics

Metrics are collected for the JBoss Server application service.

Table 1-21. JBoss Server Metrics

Metric Name	Category	KPI
Buffer Pool<InstanceName> Count	Jboss Server	False
Buffer Pool<InstanceName> Memory Used	Jboss Server	False
Buffer Pool<InstanceName> Total Capacity	Jboss Server	False
Class Loading Loaded Class Count	Jboss Server	False
Class Loading Total Loaded Class Count	Jboss Server	False
Class Loading Unloaded Class Count	Jboss Server	False
File Descriptor Usage Max File Descriptor Count	Jboss Server	False
File Descriptor Usage Open File Descriptor Count	Jboss Server	False
Http Listener<InstanceName> Bytes Received	Jboss Server	False
Http Listener<InstanceName> Bytes Sent	Jboss Server	False
Http Listener<InstanceName> Error Count	Jboss Server	False

Table 1-21. JBoss Server Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	Jboss Server	False
Garbage Collection<InstanceName> Total Collection Count	Jboss JVM Garbage Collector	False
Garbage Collection<InstanceName> Total Collection Time	Jboss JVM Garbage Collector	False
JVM Memory Heap Memory Usage Committed Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Initial Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Maximum Memory	Jboss JVM Memory	False
JVM Memory Heap Memory Usage Used Memory	Jboss JVM Memory	True
JVM Memory Non Heap Memory Usage Committed Memory	Jboss JVM Memory	False
JVM Memory Non Heap Memory Usage Initial Memory	Jboss JVM Memory	False

Table 1-21. JBoss Server Metrics (continued)

Metric Name	Category	KPI
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

HyperV Metrics

Metrics are collected for the HyperV application service.

Table 1-22. HyperV 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 O 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
Memory Cache Bytes	HyperV	False
Memory Cache Faults	HyperV	False

Table 1-22. HyperV Metrics (continued)

Metric Name	Category	KPI
Memory Pages	HyperV	False
Network<instanceName> Packets Outbound Error	HyperV	False
Network<instanceName> Packets Received Error	HyperV	False
Application Availability	HyperV	False

Oracle DB Metrics

Metrics are collected for the Oracle DB application service.

Oracle DB cannot be activated on Linux platforms.

Table 1-23. Oracle DB Metrics

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

Table 1-23. Oracle DB Metrics (continued)

Metric Name	Category	KPI
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
Utilization User Transaction Persec	OracleDB	False
Utilization Database Time Persc	OracleDB	False
Application Availability	Oracle DB	False

Cassandra Metrics

Metrics are collected for the Cassandra application service.

Table 1-24. Cassandra 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 1-24. Cassandra 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 1-24. Cassandra 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
Application Availability	Cassandra	False

MongoDB Metrics

Metrics are collected for the MongoDB application service.

Table 1-25. 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

Table 1-25. MongoDB Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	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 Metrics

Metrics are collected for the MS Exchange application service.

Table 1-26. MS Exchange 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

Table 1-26. MS Exchange Metrics (continued)

Metric Name	Category	KPI
Active Manager Server Server-Side Calls per second	MS Exchange	True
Active Manager Server Total Number of Databases	MS Exchange	True
ActiveSync Average Request Time	MS Exchange	True
ActiveSync Current Requests	MS Exchange	False
ActiveSync Mailbox Search Total	MS Exchange	False
ActiveSync Ping Commands Pending	MS Exchange	False
ActiveSync Requests per second	MS Exchange	True
ActiveSync Sync Commands per second	MS Exchange	True
ASP.NET Application Restarts	MS Exchange	False
ASP.NET Request Wait Time	MS Exchange	True
ASP.NET Worker Process Restarts	MS Exchange	False
Autodiscover Service Requests per second	MS Exchange	True
Availability Service Average Time to Process a Free Busy Request	MS Exchange	True
Outlook Web Access Average Search Time	MS Exchange	True
Outlook Web Access Requests per second	MS Exchange	False
Outlook Web Access Current Unique Users	MS Exchange	False
Application Availability	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

Table 1-26. MS Exchange Metrics (continued)

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

Microsoft SQL Server Metrics

Metrics are collected for the Microsoft SQL Server application service.

Table 1-27. Microsoft SQL Server Metrics

Metric Name	Category	KPI
CPU<InstanceName> CPU Usage (%)	Microsoft SQL Server	False
Database IO Rows Reads Bytes/Sec	Microsoft SQL Server	False

Table 1-27. Microsoft SQL Server Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-27. Microsoft SQL Server Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-27. Microsoft SQL Server Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-27. Microsoft SQL Server Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	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 1-28. 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

Table 1-28. MySQL Metrics (continued)

Metric Name	Category	KPI
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
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

Table 1-28. MySQL Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	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 1-29. 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
Application Availability	Nginx	False

Network Time Protocol Metrics

Metrics are collected for the Network Time Protocol application service.

Table 1-30. Network Time Protocol Metrics

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

Oracle WebLogic Server Metrics

Metrics are collected for the Oracle WebLogic Server application service.

Table 1-31. Oracle WebLogic Server 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
Application Availability	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

Table 1-31. Oracle WebLogic Server Metrics (continued)

Metric Name	Category	KPI
UTILIZATION Usage	Weblogic JVM Memory Pool	False
UTILIZATION UpTime	Weblogic JVM Runtime	False

Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

Table 1-32. 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

Table 1-32. Pivotal TC Server Metrics (continued)

Metric Name	Category	KPI
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
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
Application Availability	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

Table 1-32. Pivotal TC Server Metrics (continued)

Metric Name	Category	KPI
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 1-33. 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
Application Availability	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

Table 1-33. PostgreSQL (continued)

Metric Name	Category	KPI
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 1-34. 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

Table 1-34. RabbitMQ Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	RabbitMQ	False
Messages Publish in	RabbitMQ Exchange	False
Messages Publish out	RabbitMQ Exchange	False
Consumer Utilisation	RabbitMQ Queue	False
Consumers	RabbitMQ Queue	False
Memory	RabbitMQ Queue	False
Messages Ack	RabbitMQ Queue	False
Messages Ack rate	RabbitMQ Queue	False
Messages Deliver	RabbitMQ Queue	False
Messages Deliver get	RabbitMQ Queue	False
Messages Persist	RabbitMQ Queue	False
Messages Publish	RabbitMQ Queue	False
Messages Publish rate	RabbitMQ Queue	False
Messages Ram	RabbitMQ Queue	False
Messages Ready	RabbitMQ Queue	False
Messages Redeliver	RabbitMQ Queue	False
Messages Redeliver rate	RabbitMQ Queue	False
Messages Space	RabbitMQ Queue	False
Messages Unack	RabbitMQ Queue	False
Messages Unacked	RabbitMQ Queue	False
Messages	RabbitMQ Queue	False

There are no metrics collected for RabbitMQ Virtual Host.

Riak KV Metrics

Metrics are collected for the Riak KV application service.

Table 1-35. Riak KV 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
Application Availability	Riak KV	False

SharePoint Server Metrics

Metrics are collected for the SharePoint Server application service.

Table 1-36. SharePoint Server 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

Table 1-36. SharePoint Server Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	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 1-37. 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

Table 1-37. WebSphere Metrics (continued)

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

Table 1-37. WebSphere Metrics (continued)

Metric Name	Category	KPI
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
Application Availability	WebSphere	False

Windows Service Metrics

Metrics are collected for Windows services.

Table 1-38. Windows Service Metrics

Metric Name	Category	KPI
AVAILABILITY Resource Availability	Services	False
UTILIZATION Memory Usage(%)	Services	False
UTILIZATION CPU Usage(%)	Services	False

Linux Process Metrics

Metrics are collected for Linux services.

Table 1-39. Linux Process Metrics

Metric Name	Category	KPI
AVAILABILITY Resource Availability	Processes	False
UTILIZATION Memory Usage (%)	Processes	False

Table 1-39. Linux Process Metrics (continued)

Metric Name	Category	KPI
UTILIZATION CPU Usage (%)	Processes	False
UTILIZATION Number of Processes	Processes	False

Remote Check Metrics

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

HTTP Metrics

vRealize Operations discovers metrics for HTTP remote checks.

HTTP Metrics

Table 1-40. HTTP Metrics

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

ICMP Metrics

vRealize Operations discovers metrics for the ICMP object type.

Table 1-41. 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 discovers metrics for the TCP object type.

Table 1-42. TCP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	False

UDP Metrics

vRealize Operations discovers metrics for the UDP object type.

Table 1-43. UDP Metrics

Metric Name	KPI
Availability	False
Response Time	True
Result Code	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 1-44. 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

Table 1-44. VeloCloud Gateway Metrics (continued)

Component	Metrics
	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
	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 1-45. VeloCloud Orchestrator Metrics

Component	Metrics
General	General Free Memory (%)
	General Status

Metrics - Ngnix

Metrics are collected for the VeloCloud Ngnix.

Table 1-46. Ngnix 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 1-47. Redis Metrics

Component	Metrics
Publish Subscribe.	Publish Subscribe Channels
Total	Total Commands Processed
	Total Connections Received
Used	Used CPU
	Used Memory
	Used Peak Memory

Metrics - ClickHouse

Metrics are collected for the VeloCloud Clickhouse.

Table 1-48. 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)

Table 1-48. Clickhouse Metrics (continued)

Component	Metrics
Events	Buffers Storage Buffer Rows
	Events Context Lock
	Events Disk Write Elapsed (μ s)
	Events File Open
	Events Function Execute
	Events Hard Page Faults
	Events Lock Readers Wait (μ s)
	Events OS IO wait (ms)
	Events OS Write (Bytes)
	Events Query
	Events Readers Wait (ms)
	Events Real Time
	Events Soft Page Faults (μ s)
	Events System Time (μ s)
	Events User Time (μ s)
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
Summary	Replicas Total Queue Size
	Summary Background Pool Task

Table 1-48. Clickhouse Metrics (continued)

Component	Metrics
	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
	Summary Zookeeper Session
	Summary Zookeeper Watch
Write Buffer	Write Buffer File Descriptor Write
Replicated	Replicated Fetch
Memory	Memory Tracking
Query	Query Thread

Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

Table 1-49. 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.

Table 1-49. Virtual Machine Metrics (continued)

Metric Name	Description
Guest OS Services Number of OOTB Services	Number of out-of-the-box services discovered in the VM.
Guest OS Services Number of Outgoing Connections	Number of outgoing connection counts from the discovered services.
Guest OS Services Number of Incoming Connections	Number of incoming connection counts to the discovered services.

Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

Table 1-50. Service Summary Metrics

Metric Name	Description
Summary Incoming Connections Count	Number of incoming connections.
Summary Outgoing Connections Count	Number of outgoing connections.
Summary Connections Count	Number of incoming and outgoing connections.
Summary Pid	Process ID.

Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

Table 1-51. 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 1-52. Service Type Metrics

Metric Name	Description
Number of instances	Number of instances of this service type.

Calculated Metrics

vRealize Operations 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 calculates metrics for objects of type:

- vSphere World
- Virtual Machine
- Host System
- Datastore

From data that the vRealize Operations adapter collects, vRealize Operations 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 1-53. 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

Metric Name	Description
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
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

Metric Name	Description
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
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
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 cpul 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

Metric Name	Description
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 cpudemand 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 cpu demand recommendedSize
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

Metric Name	Description
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
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

Metric Name	Description
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
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

Metric Name	Description
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
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

Metric Name	Description
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 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 1-54. 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 1-55. 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 , 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
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

Table 1-55. System Metrics (continued)

Metric Name	Description
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 1-56. 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

vRealize Operations uses the vRealize Operations adapter to collect metrics that monitor its own performance. These self-monitoring metrics drive capacity models for vRealize Operations objects and are useful for diagnosing problems with vRealize Operations .

Analytics Metrics

vRealize Operations collects metrics for the vRealize Operations analytics service, including threshold checking metrics.

Table 1-57. 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.
LocalPrimaryFDAItSize	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 1-58. Overall Threshold Checking Metrics for the Analytics Service

Metric Key	Metric Name	Description
Count	Count	Count
Duration TotalDuration	Total	Total length of duration (ms)
Duration AvgDuration	Average	Average duration (ms)
Duration MinDuration	Minimum	Minimum duration (ms)
Duration MaxDuration	Maximum	Maximum duration (ms)
IncomingObservationsSize TotalCount	Total	Total
IncomingObservationsSize AvgCount	Average	Average
IncomingObservationsSize MinCount	Minimal	Minimal
IncomingObservationsSize MaxCount	Maximal	Maximal
CheckThresholdAndHealth Count	Count	Count
CheckThresholdAndHealth Duration TotalDuration	Total	Total length of duration (ms)
CheckThresholdAndHealth Duration AvgDuration	Average	Average duration (ms)
CheckThresholdAndHealth Duration MinDuration	Minimum	Minimum duration (ms)
CheckThresholdAndHealth Duration MaxDuration	Maximum	Maximum duration (ms)
CheckThresholdAndHealth OutcomeObservationsSize TotalCount	Total	Total
CheckThresholdAndHealth OutcomeObservationsSize AvgCount	Average	Average
CheckThresholdAndHealth OutcomeObservationsSize MinCount	Minimal	Minimal
CheckThresholdAndHealth OutcomeObservationsSize MaxCount	Maximal	Maximal
SuperMetricComputation Count	Count	Count
SuperMetricComputation Duration TotalDuration	Total	Total length of duration (ms)

Table 1-58. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
SuperMetricComputation Duration AvgDuration	Average	Average duration (ms)
SuperMetricComputation Duration MinDuration	Minimum	Minimum duration (ms)
SuperMetricComputation Duration MaxDuration	Maximum	Maximum duration (ms)
SuperMetricComputation SuperMetricsCount TotalCount	Total	Total
SuperMetricComputation SuperMetricsCount AvgCount	Average	Average
SuperMetricComputation SuperMetricsCount MinCount	Minimal	Minimal
SuperMetricComputation SuperMetricsCount MaxCount	Maximal	Maximal
StoreObservationToFSDB Count	Count	Count
StoreObservationToFSDB Duration TotalDuration	Total	Total length of duration (ms)
StoreObservationToFSDB Duration AvgDuration	Average	Average duration (ms)
StoreObservationToFSDB Duration MinDuration	Minimum	Minimum duration (ms)
StoreObservationToFSDB Duration MaxDuration	Maximum	Maximum duration (ms)
StoreObservationToFSDB StoredObservationsSize TotalCount	Total	Total
StoreObservationToFSDB StoredObservationsSize AvgCount	Average	Average
StoreObservationToFSDB StoredObservationsSize MinCount	Minimal	Minimal
StoreObservationToFSDB StoredObservationsSize MaxCount	Maximal	Maximal
UpdateResourceCache Count	Count	Count
UpdateResourceCache Duration TotalDuration	Total	Total
UpdateResourceCache Duration AvgDuration	Average	Average
UpdateResourceCache Duration MinDuration	Minimum	Minimum
UpdateResourceCache Duration MaxDuration	Maximum	Maximum

Table 1-58. Overall Threshold Checking Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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
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 1-59. 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)

Table 1-59. Dynamic Threshold Calculation Metrics for the Analytics Service (continued)

Metric Key	Metric Name	Description
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
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 1-60. 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 collects metrics for the vRealize Operations Collector service objects.

Table 1-61. 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.
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 collects metrics for the vRealize Operations Controller objects.

Table 1-62. 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 collects metrics for the vRealize Operations file system database (FSDB) objects.

Table 1-63. 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 1-64. Storage Thread Pool Metrics for FSDB

Metric Key	Metric Name	Description
StoreOperationsCount	Store operations count	Store operations count
StorageThreadPool Duration TotalDuration	Total	Total number of duration (ms)
StorageThreadPool Duration AvgDuration	Average	Average duration (ms)
StorageThreadPool Duration MinDuration	Minimum	Minimum duration (ms)
StorageThreadPool Duration MaxDuration	Maximum	Maximum duration (ms)
StorageThreadPool SavedMetricsCount TotalCount	Total	Total
StorageThreadPool SavedMetricsCount AvgCount	Average	Average
StorageThreadPool SavedMetricsCount MinCount	Minimal	Minimal
StorageThreadPool SavedMetricsCount MaxCount	Maximal	Maximal

Product UI Metrics

vRealize Operations collects metrics for the vRealize Operations product user interface objects.

Table 1-65. Product UI Metrics

Metric Key	Metric Name	Description
ActiveSessionsCount	Active sessions	Active sessions
CurrentHeapSize	Current heap size	Current heap size.

Table 1-65. Product UI Metrics (continued)

Metric Key	Metric Name	Description
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 1-66. API Call Metrics for the Product UI

Metric Key	Metric Name	Description
APICalls HTTPRequesterRequestCount	HTTPRequester request count	HTTPRequester request count
APICalls AvgHTTPRequesterRequestTime	HTTPRequester average request time	HTTPRequester average request time (ms)
APICalls FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
APICalls AvgAlertRequestTime	Average alert request time	Average alert request time (ms)
APICalls AlertRequestCount	Alert request count	Alert request count
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 collects metrics for the vRealize Operations administration user interface objects.

Table 1-67. 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 1-68. 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 collects metrics for the vRealize Operations API objects.

Table 1-69. 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 1-70. 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 1-71. 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 collects metrics for vRealize Operations Cluster and Slice Administration (CaSA) objects.

Table 1-72. 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 1-73. API Call Metrics for Cluster and Slice Administration

Metric Key	Metric Name	Description
API Calls TotalRequests	Total Requests	Total Requests
API Calls AvgResponseTime	Average Response Time	Average Response Time (ms)
API Calls MinResponseTime	Minimum Response Time	Minimum Response Time (ms)
API Calls MaxResponseTime	Maximum Response Time	Maximum Response Time (ms)
API Calls ServerErrorResponseCount	Server Error Response Count	Server Error Response Count
API Calls FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
API Calls FailedAuthorizationCount	Minimum Response Time	Minimum Response Time (ms)

Watchdog Metrics

vRealize Operations collects watchdog metrics to ensure that the vRealize Operations services are running and responsive.

Watchdog Metrics

The watchdog metric provides the total service count.

Table 1-74. Watchdog Metrics

Metric Key	Metric Name	Description
ServiceCount	Service Count	Service Count

Service Metrics

Service metrics provide information about watchdog activity.

Table 1-75. Metrics for the vRealize Operations 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 collects metrics for the vRealize Operations node objects.

Metrics can be calculated for node objects. See [Calculated Metrics](#).

Table 1-76. Node Metrics

Metric Key	Metric Name	Description
Component Count	Component count	The number of vRealize Operations 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 1-77. 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 1-78. 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 1-79. Resource Limit Metrics for the Node

Metric Key	Metric Name	Description
resourceLimit numProcesses	Number of processes	Number of processes
resourceLimit openFiles	Number of open files	Number of open files
resourceLimit openFilesMax	Number of open files maximum limit	Number of open files maximum limit
resourceLimit numProcessesMax	Number of processes maximum limit	Number of processes maximum limit

Table 1-80. Network Metrics for the Node

Metric Key	Metric Name	Description
net allInboundTotal	All inbound connections	All inbound total
net allOutboundTotal	All outbound connections	All outbound total
net tcpBound	TCP bound	TCP bound
net tcpClose	TCP state CLOSE	Number of connections in TCP CLOSE
net tcpCloseWait	TCP state CLOSE WAIT	Number of connections in TCP state CLOSE WAIT
net tcpClosing	TCP state CLOSING	Number of connections in TCP state CLOSING
net tcpEstablished	TCP state ESTABLISHED	Number of connections in TCP state ESTABLISHED
net tcpIdle	TCP state IDLE	Number of connections in TCP state IDLE
net tcpInboundTotal	TCP inbound connections	TCP inbound connections
net tcpOutboundTotal	TCP outbound connections	TCP outbound connections
net tcpLastAck	TCP state LAST ACK	Number of connections in TCP state LAST ACK

Table 1-80. Network Metrics for the Node (continued)

Metric Key	Metric Name	Description
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 1-81. Network Interface Metrics for the Node

Metric Key	Metric Name	Description
net iface speed	Speed	Speed (bits/sec)
net iface rxPackets	Receive packets	Number of received packets
net iface rxBytes	Receive bytes	Number of received bytes
net iface rxDropped	Receive packet drops	Number of received packets dropped
net iface rxFrame	Receive packets frame	Number of receive packets frame
net iface rxOverruns	Receive packets overruns	Number of receive packets overrun
net iface txPackets	Transmit packets	Number of transmit packets
net iface txBytes	Transmit bytes	Number of transmit bytes
net iface txDropped	Transmit packet drops	Number of transmit packets dropped
net iface txCarrier	Transmit carrier	Transmit carrier
net iface txCollisions	Transmit packet collisions	Number of transmit collisions
net iface txErrors	Transmit packet errors	Number of transmit errors
net iface txOverruns	Transmit packet overruns	Number of transmit overruns

Table 1-82. 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

Table 1-82. Disk Filesystem Metrics for the Node (continued)

Metric Key	Metric Name	Description
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 1-83. 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 1-84. 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 1-85. 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 1-86. CPU Metrics for the Node

Metric Key	Metric Name	Description
cpulcombined	Combined load	Combined load (User + Sys + Nice + Wait)
cpulidle	Idle	Idle time fraction of total available cpu (cpu load)

Table 1-86. CPU Metrics for the Node (continued)

Metric Key	Metric Name	Description
cpu irq	Irq	Interrupt time fraction of total available cpu (cpu load)
cpu nice	Nice	Nice time fraction of total available cpu (cpu load)
cpu softIrq	Soft Irq	Soft interrupt time fraction of total available cpu (cpu load)
cpu stolen	Stolen	Stolen time fraction of total available cpu (cpu load)
cpu sys	Sys	Sys time fraction of total available cpu (cpu load)
cpu user	User (cpu load)	User time fraction of total available cpu (cpu load)
cpu wait	Wait (cpu load)	Wait time fraction of total available cpu (cpu load)
cpu total	Total available for a cpu	Total available for a cpu
cpu allCpuCombined	Total combined load for all cpus	Total combined load for all cpus (cpu load)
cpu allCpuTotal_ghz	Available	Available
cpu allCpuCombined_ghz	Used	Used
cpu allCpuCombined_percent	CPU usage	CPU usage (%)

Table 1-87. 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 1-88. 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 1-89. 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 1-90. 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 collects metrics for the vRealize Operations 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 1-91. 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
Count (VM)	Usage Count	This metric displays how many units of the license capacity is currently used.
License Metrics for Cluster Object		

Table 1-91. Cluster Metrics (continued)

Metric Key	Metric Name	Description
Used (VM)	Usage (%)	<p>This metric displays the percentage of the total license capacity currently used.</p> <p>Note An alert is generated if the license threshold is in any of the following states:</p> <ul style="list-style-type: none"> ■ >= 80% - Warning ■ >= 90% - Immediate ■ =95% - Catastrophic
Days Remaining (day)	Days Remaining	This metric displays the days remaining before the license expires.
Capacity	Capacity	Displays the maximum number of units (of the given capacity type) that can be licensed by this license key.
Type	Type	Displays the license type for the cluster object.
Expiry	Expiration Date	Displays the date when the license expires.
Capacity	Capacity	Displays the maximum number of units (of the given capacity type) that can be licensed by this license key.

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 1-92. DT Metrics for the Cluster

Metric Key	Metric Name	Description
dt isRunning	Running	Running
dt dtRunTime	Running duration	Running duration (ms)
dt StartTime	Running start time	Running start time
dt percentage	Percent	Percent (%)
dt executorCount	Executor Node Count	Executor Node Count
dt resourceCount	Resource Count	Resource Count
dt fsdbReadTime	FSDB Read Time	FSDB Read Time (ms)
dt dtObjectSaveTime	DT Object Save Time	DT Object Save Time (ms)

Table 1-92. DT Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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 1-93. 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 1-94. Gemfire cluster Metrics for the Cluster

Metric Key	Metric Name	Description
GemfireCluster System AvgReads	Average reads per second	The average number of reads per second for all members
GemfireCluster System AvgWrites	Average writes per second	The average number of writes per second for all members
GemfireCluster System DiskReadsRate	Disk reads rate	The average number of disk reads per second across all distributed members
GemfireCluster System DiskWritesRate	Disk writes rate	The average number of disk writes per second across all distributed members
GemfireCluster System GarbageCollectionCount	Total garbage collection count	The total garbage collection count for all members

Table 1-94. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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)
GemfireCluster System Member FunctionExecutionRate	Function executions per second	Function executions per second

Table 1-94. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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
GemfireCluster Region CreatesRate	Creates per second	Creates per second

Table 1-94. Gemfire cluster Metrics for the Cluster (continued)

Metric Key	Metric Name	Description
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 1-95. 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 1-96. 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 1-97. 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 1-98. 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 1-99. 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 collects metrics for various persistence resources or service groups.

Activity Metrics

Activity metrics relate to the activity framework.

Table 1-100. 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 1-101. 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 1-102. 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 1-102. 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 1-103. 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 1-103. 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 collects metrics for the vRealize Operations remote collector node objects.

Table 1-104. Remote Collector Metrics

Metric Key	Metric Name	Description
ComponentCount	Component Count	The number of vRealize Operations Objects reporting for this node.

Table 1-105. 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 1-106. 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 1-107. 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 1-108. 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 1-108. 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 1-109. 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 1-109. 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 1-110. 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 1-111. 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 1-112. 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 1-113. 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 1-114. CPU Metrics for the Remote Collector

Metric Key	Metric Name	Description
cpu combined	Combined load	Combined load (User + Sys + Nice + Wait)
cpulidle	Idle	Idle time fraction of total available cpu (cpu load)
cpulirq	Irq	Interrupt time fraction of total available cpu (cpu load)
cpunice	Nice	Nice time fraction of total available cpu (cpu load)
cpu softIrq	Soft Irq	Soft interrupt time fraction of total available cpu (cpu load)
cpustolen	Stolen	Stolen time fraction of total available cpu (cpu load)
cpusys	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 1-115. 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 1-116. 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 1-117. 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 1-118. Blueprint Metrics

Property Name	Metrics
Summary	VMCount

Project Metrics

vRealize Automation 8.x collects metrics for objects such as project object.

Table 1-119. 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 1-120. 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 1-121. Organization Metrics

Property Name	Metrics
Summary	TotalBlueprints
Summary	TotalProjects
Summary	VMCount

Table 1-121. 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 1-122. 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 1-123. 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 1-124. Cloud Automation Services Entity Status Metrics

Property Name	Metrics
Summary	TotalClusters

Cloud Zone Limits Metrics

vRealize Automation 8.x collects memory, storage, and vCPU limit related metrics at project level on per cloud zone basis.

Cloud Zone Limits Metrics

You can view the memory, storage, and vCPU limit related metrics defined in vRealize Automation 8.x in vRealize Operations . vRealize Operations alerts when the utilisation exceeds configured limits in any of the individual cloud zones.

Note The resource limits are configured in vRealize Automation 8.x. For vCenter objects, limits for all resources - vCPU, memory, and storage - can be configured. However, for public clouds like AWS, Azure, and GCP only vCPU and memory limits can be configured.

Table 1-125. Cloud Zone Limits Metrics

Property Name	Metrics	Description
Cloud Zone Limits	Memory allocated (KB)	This metric displays the total utilized cloud zone memory.
	Memory Limit (KB)	This metric displays the limit set on the cloud zone memory usage.
	Memory Utilized (%)	This metric displays the percentage of memory used out of the memory limit set. Formula $((\text{Memory allocated (KB)} / \text{Memory Limit (KB)})$
	Storage allocated (GB)	This metric displays the total storage utilized by the cloud zone.
	Storage Limit (GB)	This metrics displays the limit set on the cloud zone storage usage.
	Storage Utilized (%)	This metric displays the percentage of storage used out of the storage limit set. Formula $(\text{Storage allocated (KB)} / \text{Storage Limit (KB)})$
	vCPU allocated	This metric displays the total vCPU utilized by the cloud zone.
	vCPU Limit	This metrics displays the limit set on the cloud zone vCPU usage.
	vCPU Utilized (%)	This metric displays the percentage of vCPU used out of the vCPU limit set. Formula $(\text{vCPU allocated (KB)} / \text{vCPU Limit (KB)})$

Metrics for vSAN

vRealize Operations collects metrics for vSAN objects.

From the left menu, click **Environment > Object Browser > 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 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 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 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 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 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 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 collects the metrics you use to monitor the performance of your vSAN cluster.

vRealize Operations 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 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 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 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 Drive 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	<ul style="list-style-type: none"> ■ 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 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 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 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 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 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 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 Microsoft Azure

vRealize Operations collects metrics for Microsoft Azure adapter objects.

From the left menu, click **Environment > Object Browser > 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 .

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 .

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 .

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 .

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 .

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 .

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 .

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 .

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 components.

EC2 Metrics

The following metrics are available for each EC2 instance in your vRealize Operations 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 1-126. 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 1-126. 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 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 1-127. 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 1-127. 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 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 1-128. 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 environment.

Table 1-129. 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 environment.

Table 1-130. 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 1-130. 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 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 1-131. 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 1-131. 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 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 1-132. 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 1-132. 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 1-132. 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 environment.

Table 1-133. 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 1-133. 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 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 1-134. 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 1-134. 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 1-134. ElastiCache Cache Node Metrics (continued)

Name	Category	Type	Unit	Instanced
CacheMisses	Performance	Metric	Count	No
ReplicationLag	Performance	Metric	Count	No

Note The Instanced column shows if the given metrics have any the accumulated values from the child entities or not.

RDS DB Instance Metrics

The following metrics are available for each RDS DB instance in your vRealize Operations environment.

Table 1-135. 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 environment.

Table 1-136. 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 environment.

Table 1-137. 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 environment.

Table 1-138. 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 environment.

Table 1-139. 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 1-139. 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 environment.

Table 1-140. 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 environment.

Table 1-141. 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 environment.

Table 1-142. 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 environment.

Table 1-143. 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 environment.

Table 1-144. 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 1-144. 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 environment.

Table 1-145. 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 1-145. 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 environment.

Table 1-146. 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 1-146. 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 environment.

Table 1-147. 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 1-147. 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 environment.

Table 1-148. 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 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 1-149. 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

Table 1-149. ElasticCache Cache Cluster Metrics (continued)

Name	Category	Type	Unit	Instanced
BytesUsedForCache	Memory	Metric	Bytes	No
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 environment.

Table 1-150. EFS Metrics

Service	Metrics
EFS	BurstCreditBalance
	ClientConnections
	DataReadIOBytes
	DataWriteIOBytes

Table 1-150. EFS Metrics (continued)

Service	Metrics
	MetadatalOBytes
	PercentIOLimit
	PermittedThroughput
	TotalIOBytes

Elastic Beanstalk Environment Metrics

The following metrics are available for each Elastic Beanstalk Environment instance in your vRealize Operations environment.

Table 1-151. 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

Table 1-151. Elastic Beanstalk Environment Metrics (continued)

Service	Metrics
	InstancesOk
	InstancesWarning
	EnvironmentHealth

AWS Transit Gateway Metrics

The following metrics are available for each AWS Transit Gateway instance in your vRealize Operations environment.

Table 1-152. 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 environment.

Table 1-153. 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

Table 1-153. EKS Cluster Metrics (continued)

Service	Metrics
	node_cpu_utilization
	node_filesystem_utilization
	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 1-154. 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.

Table 1-154. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	Cost Monthly Total Expense	Double	Represents the total amount spent on the OnDemand and Commit purchases for a month.
	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	

Table 1-154. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
	Configuration Maximum Max clusters Soft Limit % Used	Double	Represents the maximum number of ESXi hosts per SDDC.
	Configuration Maximum Max clusters Hard Limit % Used	Double	
	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	

Table 1-154. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
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.
	Configuration Maximum Minimum hosts per cluster for full SLA Limit Violated	Double	
	Configuration Maximum Min hosts per cluster for no SLA Limit	Double	Represents the minimum number of ESXi hosts per vSphere cluster with no SLA.
	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	

Table 1-154. VMware Cloud on AWS Metrics (continued)

Object Type	Metric Key	Metric Value	Description
Logical Switch	VMC Configuration Maximums L2VPN Client Count Limit % Used	Double	Represents the maximum number of logical segments per SDDC.
	VMC Configuration Maximums Logical Segment Count Limit	Double	
	VMC Configuration Maximums Logical Segment Count Provisioned	Double	
	VMC Configuration Maximums Logical Segment Count Limit % Used	Double	Represents the maximum number of ports on a logical segment.
	VMC Configuration Maximums Logical Ports Count Limit	Double	
	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 1-154. 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 1-154. 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 1-155. 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 1-155. 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 1-156. 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 RxFPackets ■ stats Interface TxData ■ stats Interface TxDropped ■ stats Interface TxErrors ■ stats Interface TxCarrier ■ stats Interface TxPackets ■ stats Interface TxColls

Table 1-156. 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 1-156. 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 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 1-156. 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. <hr/> <ul style="list-style-type: none"> configMax tier1RouterCount configMax routeMapCount configMax RouteMaps routeMapRuleCount <hr/> Note Metric applicable for T0 router. <hr/> <ul style="list-style-type: none"> configMax prefixListCount configMax IPPrefixLists prefixListEntriesCount <hr/> Note Metric applicable for T0 and T1 router. <hr/>
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 1-156. 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 1-156. 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 1-157. 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 1-157. 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>

Metrics in VMware vRealize Operations Management Pack for Google Cloud Platform

The VMware vRealize Operations Management Pack for Google Cloud Platform collects metrics for objects.

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics

Object Type	Metric Key
GCP World	summary Total CE Instances summary Active CE Instances summary Number of Storage Buckets summary Number of K8s Engines summary Number of BigQuery Datasets summary Number of VPN Gateways
GoogleCloudPlatformAdapterInstance	summary Total CE Instances summary Active CE Instances summary Number of Storage Buckets summary Number of K8s Engines summary Number of BigQuery Datasets summary Number of VPN Gateways

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
CE Instance	CPU utilization CPU usage_time CPU reserved_cores CPU scheduler_wait_time CPU vCPUs CPU IsSharedCpu Status Disk read_bytes_count Disk read_ops_count Disk throttled_read_bytes_count Disk throttled_read_ops_count Disk throttled_write_bytes_count Disk throttled_write_ops_count Disk write_bytes_count Disk write_ops_count Disk Disk I/O Average (Count) Disk Disk I/O Average (Bytes) Memory Total Memory maximumPersistentDisksSizeGb balloon ram_size balloon ram_used balloon swap_in_bytes_count balloon swap_out_bytes_count firewall dropped_bytes_count firewall dropped_packets_count instance_group size network received_bytes_count network received_packets_count network sent_bytes_count network sent_packets_count network Network I/O Average (Bytes) network Network I/O Average (Packets) integrity early_boot_validation_status integrity late_boot_validation_status uptime mirroring dropped_packets_count mirroring mirrored_packets_count mirroring mirrored_bytes_count

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
Storage Buckets	Summary request_count Average
	Summary request_count Maximum
	Summary request_count Minimum
	Summary request_count Sum
	Summary object_count Average
	Summary object_count Maximum
	Summary object_count Minimum
	Summary total_byte_seconds Average
	Summary total_byte_seconds Maximum
	Summary total_byte_seconds Minimum
	Summary total_bytes Average
	Summary total_bytes Maximum
	Summary total_bytes Minimum
	authz authentication_count Average
	authz authentication_count Maximum
	authz authentication_count Minimum
	authz acl_based_object_access_count Average
	authz acl_based_object_access_count Maximum
	authz acl_based_object_access_count Minimum
	authz acl_operations_count Average
	authz acl_operations_count Maximum
	authz acl_operations_count Minimum
	authz object_specific_acl_mutation_count Average
	authz object_specific_acl_mutation_count Maximum
	authz object_specific_acl_mutation_count Minimum
	network received_bytes_count Average
	network received_bytes_count Maximum
	network received_bytes_count Minimum
	network sent_bytes_count Average
	network sent_bytes_count Maximum
	network sent_bytes_count Minimum

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
K8s Clusters	CPU utilization Average CPU utilization Maximum CPU utilization Minimum CPU usage_time CPU reserved_cores disk bytes_total Average disk bytes_used Average disk bytes_used Maximum disk bytes_used Minimum disk inodes_free disk inodes_total memory bytes_total Average memory bytes_used Average memory bytes_used Maximum memory bytes_used Minimum memory page_fault_count uptime pid_used pid_limit accelerator memory_total Average accelerator memory_used Average accelerator memory_used Maximum accelerator memory_used Minimum accelerator memory_used accelerator duty_cycle accelerator request

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
Big Query DataSet	slots allocated slots allocated_for_project slots allocated_for_project_and_job_type slots allocated_for_reservation slots total_allocated_for_reservation slots total_available query count query execution_times query scanned_bytes query scanned_bytes_billed query statement_scanned_bytes query statement_scanned_bytes_billed job num_in_flight storage table_count configuration dataSetId configuration generatedId configuration creationTime configuration lastModifiedTime configuration defaultTableExpiration configuration description configuration etag configuration location configuration friendlyName configuration labels
Big Query Tables	storage stored_bytes storage uploaded_bytes storage uploaded_bytes_billed storage uploaded_row_count configuration tableId configuration generatedId configuration creationTime configuration lastModifiedTime configuration expirationTime configuration description configuration etag configuration location configuration friendlyName configuration labels configuration numBytes configuration numLongTermBytes configuration numRows configuration datasourceUri configuration datasourceFormat

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
VPN Gateways	gateway connections
VPN Tunnels	network dropped_received_packets_count network dropped_sent_packets_count network received_bytes_count network received_packets_count network sent_bytes_count network sent_packets_count tunnel_established

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
K8s Nodes	CPU utilization CPU usage_time CPU reserved_cores CPU scheduler_wait_time CPU vCPUs CPU IsSharedCpu Status Disk read_bytes_count Disk read_ops_count Disk throttled_read_bytes_count Disk throttled_read_ops_count Disk throttled_write_bytes_count Disk throttled_write_ops_count Disk write_bytes_count Disk write_ops_count Disk Disk I/O Average (Count) Disk Disk I/O Average (Bytes) Memory Total Memory maximumPersistentDisksSizeGb balloon ram_size balloon ram_used balloon swap_in_bytes_count balloon swap_out_bytes_count firewall dropped_bytes_count firewall dropped_packets_count instance_group size network received_bytes_count network received_packets_count network sent_bytes_count network sent_packets_count network Network I/O Average (Bytes) network Network I/O Average (Packets) integrity early_boot_validation_status integrity late_boot_validation_status uptime mirroring dropped_packets_count mirroring mirrored_packets_count mirroring mirrored_bytes_count
K8s Pods	network received_bytes_count network sent_bytes_count volume total_bytes volume used_bytes volume utilization

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
K8s Container	CPU core_usage_time CPU limit_cores CPU limit_utilization CPU request_cores CPU request_utilization memory limit_bytes memory limit_utilization memory page_fault_count memory request_bytes memory request_utilization memory used_bytes memory restart_count uptime
Region	summary Total CE Instances summary Active CE Instances summary Number of Storage Buckets summary Number of K8s Engines summary Number of BigQuery Datasets
Cloud SQL	cpu utilization disk bytes_used disk write_ops_count disk read_ops_count memory usage network sent_bytes_count network received_bytes_count network connections
Memorystore Redis	stats usage memory usage stats maxmemory memory maxmemory stats total connections total stats cpu_utilization
Memorystore Memcached	node operation_count node active_connections node eviction_count node received_bytes_count node sent_bytes_count node items node hit_ratio memory utilization cpu utilization

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
Memcached Node	node operation_count node active_connections node eviction_count node received_bytes_count node sent_bytes_count node items node hit_ratio memory utilization cpu utilization
Cloud Spanner	storage limit_bytes storage utilization cpu utilization storage used_bytes api sent_bytes_count api received_bytes_count spannerDatabase (instanced metric)
Cloud Bigtable	general nodeCount cluster cpu_load cluster cpu_load_hottest_node cluster storage_utilization cluster disk_load disk bytes_used disk storage_capacity
Cloud Bigtable Cluster	general nodeCount cluster cpu_load cluster cpu_load_hottest_node cluster storage_utilization cluster disk_load disk bytes_used disk storage_capacity
Cloud Bigtable Table	server multi_cluster_failovers_count server request_count server sent_bytes_count server received_bytes_count table bytes_used

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
Firebase Realtime Database	storage total_bytes
	storage limit
	network active_connections
	network sent_bytes_count
	network monthly_sent_limit
	network monthly_sent
	Disk IO utilization
	Disk IO database_load
	Disk IO persisted_bytes_count

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
Firestore Database	Firestore in Native Mode: document read_count document write_count document delete_count network snapshot_listeners network active_connections Firestore in Datastore mode: entity read_sizes entity write_sizes api request_count index write_count
VPC Network	Quota <ul style="list-style-type: none"> ■ instances_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ internal_lb_forwarding_rules_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ internal_managed_forwarding_rules_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ internal_protocol_forwarding_rules_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ ip_aliases_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ psc_google_apis_forwarding_rules_per_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit ■ exceeded ■ psc_ilb_consumer_forwarding_rules_per_producer_vpc_network <ul style="list-style-type: none"> ■ usage ■ limit

Table 1-158. VMware vRealize Operations Management Pack for Google Cloud Platform Metrics (continued)

Object Type	Metric Key
	<ul style="list-style-type: none"> ■ exceeded ■ subnet_ranges_per_vpc_network ■ usage ■ limit ■ exceeded

Sustainability Metrics

Sustainability metrics are collected for virtual machine, host system, cluster compute resource, and vSphere World object types.

Metric Names

Metric Name	Object Type	Description
Power Total Energy (Wh)	Virtual Machine	Total energy used. Formula: Total Energy (Wh) = Sum(Power Energy (Joule))/3600
Power Total Energy (Wh)	Host System	Total energy used. Formula: Total Energy (Wh) = Sum(Power Energy (Joule))/3600
Sustainability CO2 Emission (Kg)	Cluster Compute Resource	Carbon dioxide emissions. Calculated as power consumption* CO2 Emission rate Formula: CO2 Emission (Kg) = Sum(Host System(Power Total Energy(Wh)))/1000 * 0.709
Sustainability CO2 Emission before Virtualization (Kg)	Cluster Compute Resource	Carbon dioxide emissions before virtualization, assuming that power consumption per physical server is 100W, reflecting a low end hardware specification. Formula: CO2 Emission before Virtualization (Kg) = Summary Number of Running VMs * 0.1 * 0.709

Metric Name	Object Type	Description
Sustainability CO2 Emission by Idle VMs (Kg)	Cluster Compute Resource	<p>Total carbon dioxide emission from all idle VMs. Calculated as CO2 emission rate * Power consumed by Idle VMs, where the rate is set at cluster property / 1000.</p> <p>Formula:</p> <p>CO2 Emission by Idle VMs (Kg) = Sum(VM(Power Total Energy (Wh)), If VM(Summary Reclaimable Idle = 1) * 0.709</p> <p>OR</p> <p>CO2 Emission by Idle VMs (Kg) = Power Wasted by Idle VMs (Wh) * 0.709</p>
Sustainability Electricity Cost Savings	Cluster Compute Resource	<p>Estimated cost savings by virtualizing workloads. Calculated from the difference between power consumption before virtualization and after virtualization. The electricity cost is defined at the cluster custom property.</p> <p>Formula:</p> <p>Electricity Cost Savings = (Summary Number of Running VMs * 0.1 - Sum(Host System(Power Total Energy(Wh)))/1000) * 0.108</p> <p>OR</p> <p>Daily Electricity Cost Savings = (Summary Number of Running VMs * 0.1 - Power usage (KWh)) * 0.108</p>
Sustainability Power usage (KWh)	Cluster Compute Resource	<p>Power usage calculated from all hosts in KWh.</p> <p>Formula:</p> <p>Power usage (KWh) = Sum(Host System(Power Total Energy(Wh)))/1000</p>
Sustainability Power usage per GHz (Wh)	Cluster Compute Resource	<p>Power usage efficiency. Calculated as power consumption over total GHz.</p> <p>Formula:</p> <p>Power usage per GHz (Wh) = Sum(Host System(Power Total Energy(Wh)))/CPU Usage (MHz)/1000</p>
Sustainability Power Wasted by Idle VMs (Wh)	Cluster Compute Resource	<p>Sum of electricity power used by all VMs classified as idle by the system.</p> <p>Formula:</p> <p>Power Wasted by Idle VMs (Wh) = Sum(VM(Power Total Energy (Wh)), If VM(Summary Reclaimable Idle = 1)</p>

Metric Name	Object Type	Description
Sustainability Trees to Offset Idle VMs CO2 Emission	Cluster Compute Resource	<p>Number of standard trees required to compensate CO2 emission of all Idle VMs. Based on 36.4 pound of carbon per tree.</p> <p>Formula:</p> <p>Trees to Offset Idle VMs CO2 Emission = $\text{Sum}(\text{VM}(\text{Power} \text{Total Energy (Wh)}), \text{If VM}(\text{Summary} \text{Reclaimable Idle} = 1)/1000 * 0.709 / 16.511$</p> <p>OR</p> <p>Trees to Offset Idle VMs CO2 Emission = $\text{Power Wasted by Idle VMs (Wh)} / 1000 * 0.709 / 16.511$</p>
Sustainability CO2 Emission (Kg)	vSphere World	<p>Total carbon dioxide emissions. Calculated as sum of carbon emission from all clusters.</p> <p>Formula:</p> <p>CO2 Emission (Kg) = $\text{Sum}(\text{Cluster}(\text{Sustainability} \text{CO2 Emission (Kg)}))$</p>
Sustainability CO2 Emission before Virtualization (Kg)	vSphere World	<p>Total carbon dioxide emissions before virtualization. Calculated as the sum of CO2 emission from all clusters.</p> <p>Formula:</p> <p>CO2 Emission before Virtualization (Kg) = $\text{Sum}(\text{Cluster}(\text{Sustainability} \text{CO2 Emission before Virtualization (Kg)}))$</p>
Sustainability CO2 Emission Avoided (T)	vSphere World	<p>Carbon emissions avoided with virtualization. Calculated by difference in values of carbon emissions before and after virtualization and converting value in Kg to Tonnes.</p> <p>Formula:</p> <p>CO2 Emission Avoided (T) = $(\text{CO2 Emission before Virtualization (Kg)} - \text{CO2 Emission (Kg)}) / 1000$</p>
Sustainability Electricity Cost Savings	vSphere World	<p>Total estimated cost savings by virtualizing workloads at vSphere World. Calculated as the sum of Electricity cost savings.</p> <p>Formula:</p> <p>Electricity Cost Savings = $\text{Sum}(\text{Cluster}(\text{Daily Electricity Cost Savings}))$</p>

Metric Name	Object Type	Description
Sustainability Power Savings with Virtualization (%)	vSphere World	<p>Percentage of power savings achieved by virtualization. Calculated by Formula = (Power usage before Virtualization - Power usage after Virtualization)/Power usage before Virtualization *100.</p> <p>Formula:</p> <p>Power Savings with Virtualization (%) = (Power usage Before Virtualization (KWh) - Power usage (KWh)) / Power usage Before Virtualization (KWh) * 100</p>
Sustainability Power usage (KWh)	vSphere World	<p>Power usage calculated from all hosts in KWh.</p> <p>Formula:</p> <p>Power usage (KWh) = Sum(Host System(Power Total Energy(Wh)))/ 1000</p>
Sustainability Power usage Before Virtualization (KWh)	vSphere World	<p>Power usage assuming that each low range server consumes 0.1 KWh.</p> <p>Formula:</p> <p>Power usage Before Virtualization (KWh) = Summary Number of Running VMs * 0.1</p>
Sustainability Power Wasted by Idle VMs (Wh)	vSphere World	<p>Sum of electricity power used by all VMs classified as idle by the system.</p> <p>Formula:</p> <p>Power Wasted by Idle VMs (Wh) = Sum((VM(Power Total Energy (Wh))), If VM(Summary Reclaimable Idle = 1))</p>

Property Definitions in vRealize Operations

2

Properties are attributes of objects in the vRealize Operations environment. You use properties in symptom definitions. You can also use properties in dashboards, views, and reports.

vRealize Operations 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 the *vRealize Operations Manager User Guide*.

vRealize Operations 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.

This chapter includes the following topics:

- [Properties for vCenter Server Components](#)
- [Self-Monitoring Properties for vRealize Operations](#)
- [OS and Application Monitoring Properties](#)
- [Service Discovery Properties](#)
- [Properties for vSAN](#)
- [Properties for vRealize Automation 8.x](#)
- [Properties in the NSX-T Adapter](#)
- [Placement Group Properties](#)
- [Properties for VeloCloud Gateway](#)
- [Properties for VeloCloud Orchestrator](#)
- [Sustainability Properties](#)

Properties for vCenter Server Components

The VMware vSphere solution is installed with vRealize Operations and includes the vCenter adapter. vRealize Operations 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 collects summary and event properties for vCenter Server system objects.

Table 2-1. 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 2-2. 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 2-3. 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 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 2-4. vRealize Automation Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
vRealize Automation Blueprint Name	Blueprint Name	Virtual machines deployed by vRealize Automation to be excluded from workload placements.

Table 2-5. 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 2-6. 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 2-7. 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 2-8. 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/

Table 2-8. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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)
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)

Table 2-8. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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_memfsfss	Disable unexposed features - memfsfss (isolation.tools.memSchedFakeSampleStats.disable)	Disable unexposed features - memfsfss (isolation.tools.memSchedFakeSampleStats.disable)
config security disable_unexposed_features_protocolhandler	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)
config security disable_unexposed_features_shellaction	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)
config security disable_unexposed_features_toporequest	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)
config security disable_unexposed_features_trashfolderstate	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)

Table 2-8. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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)
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)

Table 2-8. Configuration Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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 2-9. Runtime Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
runtime memoryCap	Memory Capacity	Memory Capacity

Table 2-10. CPU Usage Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
cpu limit	CPU limit	CPU limit
cpu reservation	CPU reservation	CPU reservation
cpu speed	CPU	CPU Speed

Table 2-11. 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 2-12. 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 2-13. Summary Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
summary customTag customTagValue	Value	Custom Tag Value
summary tag	vSphere Tag	vSphere Tag Name
summary parentCluster	Parent Cluster	Parent Cluster
summary parentHost	Parent Host	Parent Host
summary parentDatacenter	Parent data center	Parent data center
summary parentVcenter	Parent vCenter	Parent vCenter

Table 2-13. Summary Properties Collected for Virtual Machine Objects (continued)

Property Key	Property Name	Description
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	Appliance	Appliance
summary config productName	Product Name	Product Name
summary UUID	UUID	Unique UUID instance in vCenter that identifies all the virtual machine instances.
summary smbiosUUID	SMBIOS UUID	System Management BIOS UUID of a virtual machine.

Table 2-14. 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 2-15. 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 . This means that they do not collect data by default.

Host System Properties

vRealize Operations collects configuration, hardware, runtime, CPU, network I/O, and properties about summary use for host system objects.

Table 2-16. 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 2-16. 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 2-17. Cost Properties Collected for Host System Objects

Property Key	Property Name	Description
Cost Energy Consumed (Joule)	Energy Consumed (Joule)	Displays the energy consumed in Joules.
Cost Number of Rack Units	Number of Rack Units	Displays the number of rack units in the host.
Cost OS Categories	OS Categories	Displays the operating system categories in the host.
Cost IsServerLeased	Is Server Leased	Displays whether the server is leased or not.
Cost RemainingDepreciationMonths	Remaining Depreciation Months	Displays the remaining number of depreciation months.
Cost ServerPurchaseCost	Server Purchase Cost	Server Purchase Cost is displayed in the currency format chosen.
Cost ServerPurchaseDate	Server Purchase Date	Server Purchase Date is displayed

Table 2-18. 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 2-19. 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 2-20. 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 2-21. 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 2-22. 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 2-23. 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 2-24. 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 2-25. 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 . This means that they do not collect data by default.

Table 2-26. 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 collects configuration and summary properties for cluster compute resource objects.

Table 2-27. License Properties for Cluster Objects

Property Key	Property Name	Description
License type	License Type	Displays the license type for the cluster object
Expiry	Expiry	Displays the number of days remaining before the license expires the cluster object Note An alert is generated if the license threshold is <ul style="list-style-type: none"> ■ >= 80% - Warning ■ >= 90% - Immediate ■ =95% - Catastrophic

Table 2-28. Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
config name	Name	Name

Table 2-29. 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 2-30. 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
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 collects configuration, CPU, memory, and summary properties for resource pool objects.

Table 2-31. 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

Table 2-31. Configuration Properties Collected for Resource Pool Objects (continued)

Property Key	Property Name	Description
config memoryAllocation expandableReservation	Expandable Reservation	Memory expandable reservation
config memoryAllocation shares shares	Shares	Memory shares

Table 2-32. CPU Usage Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
cpulimit	CPU Limit	CPU Limit
cpu reservation	CPU reservation	CPU Reservation
cpu expandable_reservation	CPU expandable reservation	CPU Expandable Reservation
cpu shares	CPU Shares	CPU Shares
cpu corecount_provisioned	Provisioned vCPU(s)	Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU.

Table 2-33. 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 2-34. 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 collects configuration and summary properties for data center objects.

Table 2-35. Configuration Properties Collected for Data Center Objects

Property Key	Property Name	Description
config name	Name	Name

Table 2-36. 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 collects configuration and summary properties for storage pod objects.

Table 2-37. Configuration Properties Collected for Storage Pod Objects

Property Key	Property Name	Description
config name	Name	Name
config sdrsconfig vmStorageAntiAffinityRules	VM storage antiaffinity rules	Storage Distributed Resource Scheduler (SDRS) VM anti-affinity rules
config sdrsconfig vmdkAntiAffinityRules	VMDK antiaffinity rules	Storage Distributed Resource Scheduler (SDRS) Virtual Machine Disk (VMDK) anti-affinity rules

VMware Distributed Virtual Switch Properties

vRealize Operations collects configuration and summary properties for VMware distributed virtual switch objects.

Table 2-38. Configuration Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
config name	Name	Name

Table 2-39. 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 collects configuration and summary properties for distributed virtual port group objects.

Table 2-40. 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 2-41. 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 collects configuration, summary, and properties about datastore use for datastore objects.

Table 2-42. Capacity Properties Collected for vSAN Datastore Objects

Property Key	Property Name	Description
Capacity Available Space (GB)	Available Space	Displays the available disk space in GB.
Capacity Provisioned (GB)	Provisioned (GB)	Displays the provisioned datastore size in GB.
Capacity Total Capacity (GB)	Total Capacity (GB)	Displays the total datastore capacity in GB.
Capacity Total Provisioned Consumer Space (GB)	Total Provisioned Consumer Space (GB)	Displays the total provisioned consumer space in GB.
Capacity Used Space (GB)	Used Space (GB)	Displays the used disk space in GB.
Capacity Used Space (%)	Used Space (%)	Displays the used disk space in percentage.
Capacity Usable Capacity (GB)	Usable Capacity (GB)	Displays the usable disk capacity in GB.

Note Earlier the vSAN Datastore base rate was calculated on the basis of Total Capacity of the disk, now the vSAN datastore base rate is calculated based on the usable capacity.

Table 2-43. Summary Properties Collected for Datastore Objects

Property Key	Property Name	Description
summary vmfs_version	VMFS (Virtual Machine File System) Version	Displays the VMFS version number, contains both major version and minor version number. Note The VMFS version property is visible, only when the datastore type is VMFS.
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 2-44. 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 2-44. 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 2-45. 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 2-45. 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 . This means that they do not collect data by default.

vSphere Pod Properties

vRealize Operations collects summary and event properties for vSphere Pods.

Table 2-46. 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 2-46. 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	Memory Hot Add Configuration.
config memoryAllocation shares shares	Configuration Memory Resource Allocation Shares Shares	
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	vCPU Hot Remove Configuration.
config extraConfig mem_tps_share	Configuration Extra Configuration VM MEM TPS	N/A
config security disable_autoinstall	Configuration Security Disable tools auto install (isolation.tools.autoInstall.disable)	
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 2-46. 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 2-46. 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 2-46. 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 2-46. 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 2-46. 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 2-46. 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 2-46. 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 idle	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.
summary smbiosUUID	SMBIOS UUID	System Management BIOS UUID of a virtual machine. Note The SMBIOS UUID metric for vSphere Pod is disabled by default. You have to enable the metric at the policy level.
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.

Table 2-46. Summary Properties Collected for vSphere Pod Objects (continued)

Property Key	Localized Name	Description
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.
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 collects summary and event properties for Namespace.

Table 2-47. 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

Table 2-47. Summary Properties Collected for Namespace Objects (continued)

Property Key	Localized Name	Description
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
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 collects summary and event properties for Tanzu Kubernetes clusters.

Table 2-48. 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

vRealize Operations uses the vRealize Operations adapter to collect properties that monitor its own objects. These self-monitoring properties are useful for monitoring changes within vRealize Operations .

Analytics Properties

vRealize Operations collects properties for the vRealize Operations analytics service.

Table 2-49. 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 collects properties for the vRealize Operations node objects.

Table 2-50. 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 2-51. Memory Properties Collected for Node Objects

Property Key	Property Name	Description
mem RAM	System RAM	System RAM

Table 2-52. Service Properties Collected for Node Objects

Property Key	Property Name	Description
service proclpid	Process ID	Process ID

Remote Collector Properties

vRealize Operations collects properties for the vRealize Operations remote collector objects.

Table 2-53. 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 2-54. Memory Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
mem RAM	System RAM	System RAM

Table 2-55. Service Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
service proclpid	Process ID	Process ID

OS and Application Monitoring Properties

Properties are collected for operating systems, application services, remote checks, Linux processes, and Windows services which can be used to create reports, views, and dashboards.

Guest Information Properties

vRealize Operations displays the following guest information properties for all objects created by the OS and Application Monitoring management pack.

- Guest Info
 - Hostname
 - IP
 - OS Name
 - OS Version
 - Telegraf Version

Other properties of operating systems and application services are available under **Properties > Tags**.

Service Discovery Properties

vRealize Operations displays object properties for service discovery.

Service Discovery Adapter Instance Properties

vRealize Operations displays the following properties for the service discovery adapter instance.

Table 2-56. 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 displays the following properties for virtual machines.

Table 2-57. 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 displays the following properties for services.

Table 2-58. 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.

Table 2-58. Services Properties (continued)

Property Name	Description
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 <code>Incoming</code> . If not, it is set to <code>Outgoing</code> . If there is no connection to another service, then the value of the property is set to <code>N/A</code> .
Has Dynamic Port	Indicates whether the service has dynamic ports or not.
Status	Indicates the status of the service. Up: The service is running. Down: The service is unavailable on the monitored VM. Unavailable: The service is unavailable on a VM that is not being monitored. None: The service is not available within 7 days.

Properties for vSAN

vRealize Operations displays object properties for vSAN.

Properties for vSAN Disk Groups

vRealize Operations 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 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.

Property Name	Description
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 resolved by the DNS servers.
Configuration vSAN File Service Domain:<domainName> Gateway	Indicated 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 displays the following property for vSAN enabled host.

- Configuration|vSAN Enabled
- Configuration|vSAN|Encryption

Properties for vSAN Cache Disk

vRealize Operations 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 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 displays the following properties for vSAN file server.

- Configuration | vSAN | Primary
- Configuration | vSAN | FQDN

Properties for vSAN File Share

The vRealize Operations 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 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 displays the following properties for the NSX-T adapter.

Table 2-59. 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 2-59. 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 2-59. 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 2-59. 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 2-59. 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 environment.

Table 2-60. Placement Group Properties

Service	Property
Placement Group	State
	Strategy

Properties for VeloCloud Gateway

vRealize Operations 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 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

Sustainability Properties

In the computed metrics, constant values for carbon dioxide emission and cost of power and tree offset for carbon dioxide emission have been used.

The values are taken from the following links:

- CO2 emission per KWh = 0.709 Kg. Reference values from [Greenhouse Gas Equivalencies Calculator](#)
- Cost of Power = \$0.108 per KWh. Based on the contiguous US average value. Reference from [VMware TCO Reference Calculator](#) .
- Tree offset for CO2 Emission = 16.511 Kg (36.4 pounds of carbon per tree). Refer to [Greenhouse Gases Equivalencies Calculator](#) which is equivalent to 36.4/2.2046 Kg of carbon per tree.

- Power consumption of a small server (1 socket, 10 cores, 32 GB RAM) = 0.1 KW (Assumption).

Table 2-61.

Property Name	Object Type	Type	Default Value
Electricity Rate	Cluster Compute Resource	Numeric	0.108
CO2 Emission	Cluster Compute Resource	Numeric	0.709
Trees to Offer	Cluster Compute Resource	Numeric	16.511

Alert Definitions in vRealize Operations Manager

3

Alert definitions are a combination of symptoms and recommendations that identify problem areas in vRealize Operations 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 the *vRealize Operations Manager User Guide*.

- **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 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.

- [Alerts in VMware vRealize Operations Management Pack for Google Cloud Platform](#)

The following alerts are triggered when any of the monitoring resources in Google Cloud Platform display an unexpected behavior.

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.	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 ■ 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.	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 ■ 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.	Symptoms include the following: <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.	Symptoms include the following: <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 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.	<p>Symptoms include the following:</p> <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
IPMI System Event Log for the host is becoming full.	<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.
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 -1</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 triggers an alert with a Warning risk level.
- When 50%-75% of the child objects have risk issues, the parent object triggers an alert with an Immediate risk level.
- When 75%-100% of the child objects have risk issues, the parent object triggers an alert with a Critical risk level.

Cloud Zone Limits

- When the Project approaches 70% of Memory allocation limit in one of the cloud zones, vRealize Operations Manager triggers an alert.
- When the Project approaches 70% of Storage allocation limit in one of the cloud zones, vRealize Operations Manager triggers an alert.
- When the Project approaches 70% of vCPU allocation limits in one of the cloud zones, vRealize Operations Manager triggers an alert.

Cloud Zone

- Cloud zone has 60 days remaining until capacity runs out.
- Cloud Zone has less than 30 percent of capacity remaining.
- Cloud Zone has more than 20 percent of reclaimable capacity.

Project

- Project has more than 20 percent of reclaimable capacity.
- Project is approaching 70% of allocation limits.

vSAN Alert Definitions

vRealize Operations 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 3-1. 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.
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.

Table 3-1. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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.
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.

Table 3-1. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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.
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.

Table 3-1. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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.
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.

Table 3-1. vSAN Cluster Object Health Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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 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 3-2. 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.

Table 3-2. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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.
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 vmknic for the vSAN iSCSI target service, and verifies that all the existing targets have valid vmknic 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.

Table 3-2. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
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.
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.

Table 3-2. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vCenter server state was not pushed to ESXi due to vCenter server being out of sync.	Storage	Configuration	<p>Triggered when the vCenter server state was not pushed to ESXi due to vCenter server being out of sync.</p> <p>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.</p>
vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver.	Storage	Configuration	<p>Triggered when the vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver.</p>
vSAN build recommendation based on the available releases and VCG compatibility guide.	Storage	Availability	<p>Triggered when the vSAN build is not compatible with available releases and VCG compatibility guide.</p> <p>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.</p>
vSAN build recommendation engine has all its dependencies met and is functioning as expected.	Storage	Availability	<p>Triggered when the vSAN build recommendation engine has issues.</p> <p>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.</p>
Some disk(s) free space in vSAN Cluster is less than 10%	Storage	Capacity	<p>Triggered when the disk usage in a vSAN cluster reaches 90% of capacity.</p> <p>Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster.</p>

Table 3-2. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
Some disk(s) free space in vSAN Cluster is less than 30%	Storage	Capacity	Triggered when the disk usage in a vSAN cluster reaches 70% 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.
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.

Table 3-2. vSAN Cluster Object Risk Alert Definitions (continued)

Alert	Alert Type	Alert Subtype	Description
vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.	Storage	Performance	<p>Triggered when the vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.</p> <p>This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.</p>
vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks.	Storage	Performance	<p>Triggered when the vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks.</p> <p>This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.</p>
vSAN is using a physical disk which has high congestion value.	Storage	Performance	<p>Triggered when the vSAN is using a physical disk which has high congestion value.</p> <p>This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.</p>
vSAN iSCSI target service home object has issues.	Storage	Availability	<p>Triggered when the vSAN iSCSI target service home object has issues.</p> <p>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.</p>
vSAN iSCSI target service is not running properly or is not correctly enabled on the host.	Storage	Availability	<p>Triggered when the vSAN iSCSI target service is not running properly or is not correctly enabled on the host.</p> <p>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.</p>
vSAN performance service statistics database object is reporting issues.	Storage	Availability	<p>Triggered when the vSAN performance service statistics database object is reporting issues.</p>
vSphere cluster members do not match vSAN cluster members.	Storage	Configuration	<p>Triggered when the vSphere cluster members do not match vSAN cluster members.</p>

Table 3-3. 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.
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 respond 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.
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 AND ■ [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.

Alert Definition	Symptom	Recommendations
Virtual machine has disk I/O read latency problem.	Symptoms include all of the following: <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.	Symptoms include all of the following: <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.	Symptoms include all of the following: <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.
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.

Alert Definition	Symptom	Recommendations
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.
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.

Alert Definition	Symptom	Recommendations
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.	<p>Symptoms include all of the following:</p> <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	<p>Symptoms include all of the following:</p> <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 <p>And</p> <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.

Alerts in VMware vRealize Operations Management Pack for Google Cloud Platform

The following alerts are triggered when any of the monitoring resources in Google Cloud Platform display an unexpected behavior.

Table 3-4. Alerts in the Management Pack for Google Cloud Platform

Alert	Description
GCP CE Instance CPU Utilization is very high	Triggered with the Google Cloud Platform CE Instance CPU Utilization is very high.
GCP CE Instance is not running	Triggered with the Google Cloud Platform CE Instance is not running.
GCP K8s Container CPU usage time is very high	Triggered with the Google Cloud Platform K8s Container CPU usage time is very high.
GCP K8s Container restart count is high	Triggered with the Google Cloud Platform K8s Container restart count is high.
GCP K8s Node CPU Utilization is very high	Triggered with the Google Cloud Platform K8s Node CPU Utilization is very high
GCP K8s Node state is not running	Triggered with the Google Cloud Platform K8s Node state is not running.