vRealize Operations Definitions for Metrics, Properties, and Alerts

vRealize Operations Manager 6.6

vmware®

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

https://docs.vmware.com/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

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

Contents

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

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

Intended Audience

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

VMware Technical Publications Glossary

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

Metric Definitions in vRealize Operations Manager 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 Manager to display results that help you to manage your environment.

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

All metric definitions are provided. The metrics reported on your system depend on the objects in your environment. You can use metrics to help troubleshoot problems. See the *vRealize Operations Manager User Guide*.

Changes in Metric Availability

The CPU Demand of Recommended (%) metric is no longer available in vRealize Operations Manager version 6.x. To approximate the metric, create a super metric using the following calculations, and add it to your Views and Reports as needed.

```
( (CPU|Stress Free Demand (MHz)) × (CPU|Current Size in Unit(s)) ) ÷ ( (CPU|Recommended Size (vCPUs)) × (CPU|Current Size (MHz)) )
```

For more information about super metrics, see the vRealize Operations Manager Information Center.

This chapter includes the following topics:

Metrics for vCenter Server Components

vRealize Operations Manager connects to VMware vCenter Server® instances through the vCenter adapter to collect metrics for vCenter Server components and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

vCenter Server components are listed in the `describe.xml` file for the vCenter adapter. The following example shows sensor metrics for the host system in the `describe.xml` file.

```
<ResourceGroup instanced="false" key="Sensor" nameKey="1350" validation="">
    <ResourceGroup instanced="false" key="fan" nameKey="1351" validation="">
        <ResourceAttribute key="currentValue" nameKey="1360" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal=""  unit="percent"/>
        <ResourceAttribute key="healthState" nameKey="1361" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal="" />
    </ResourceGroup>
    <ResourceGroup instanced="false" key="temperature" nameKey="1352" validation="">
        <ResourceAttribute key="currentValue" nameKey="1362" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal="" />
        <ResourceAttribute key="healthState" nameKey="1363" dashboardOrder="1" dataType="float"
defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal="" minVal="" />
    </ResourceGroup>
 </ResourceGroup>
```

Each `ResourceAttribute` element includes the name of a metric that appears in the UI and is documented as a Metric Key.

Table 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 Manager collects CPU use, disk, memory, network, and summary metrics for objects in the vSphere world.

Capacity metrics can be calculated for vSphere world objects. See "Capacity and Project-Based Metrics," on page 77.

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Table 1-2. CPU Usage Metrics

Metric Name	Description
CPU|Capacity usage	CPU usages as a percent during the interval. Key: cpu|capacity_usagepct_average
CPU|CPU contention(%)	This metric shows the percentage of time the VMs in the ESXi hosts are unable to run because they are contending for access to the physical CPUs. The number shown is the average number for all VMs. The number will be lower than the highest number experienced by the VM that is most impacted by CPU contention. Use this metric to verify if the host can serve all its VMs efficiently. Low contention means that the VM can access everything it demands to run smoothly. It means that the infrastructure is providing good service to the application team. When using this metric, ensure that the number is within your expectation. Look at both the relative number and the absolute number. Relative means a drastic change in value, meaning that the ESXi is unable to serve the VMs. Absolute means that the real value itself is high. Investigate why the number is high. One factor that impacts this metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz, the reduction in speed is accounted for because it shows that the VM is not running at full speed. This metric is calculated in the following way: cpu|capacity_contention / (200 * summary|number_running_vcpus) Key: cpu|capacity_contentionPct
CPU|Demand (%)	This metric shows the amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. This metric represents the average active CPU load for the past five minutes. Keep this number below 100% if you set the power management to maximum. This metric is calculated in the following way: (cpu.demandmhz / cpu.capacity_provisioned)*100 Key: cpu|demandPct
CPU|Demand (MHz)	This metric shows the amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. Key: cpu|demandmhz
CPU|Demand	CPU demand in megahertz. Key: cpu|demand_average
CPU|IO wait	IO wait (ms). Key: cpu|iowait
CPU|number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Provisioned Capacity (MHz)	capacity in MHz of the physical CPU cores. Key: cpu|capacity_provisioned
CPU|Provisioned vCPU(s)	Number of provisioned CPU cores. Key: cpu|corecount_provisioned
CPU|Reserved Capacity (MHz)	Total CPU capacity reserved by virtual machines. Key: cpu|reservedCapacity_average

Table 1-2. CPU Usage Metrics (Continued)

Metric Name	Description
CPU|Usage (MHz)	CPU usages, as measured in megahertz, during the interval. ■ VM - Amount of actively used virtual CPU. This is the host's view of the CPU usage, not the guest operating system view. ■ Host - Sum of the actively used CPU of all powered on virtual machines on a host. The maximum possible value is the frequency of the two processors multiplied by the number of processors. For example, if you have a host with four 2 GHz CPUs running a virtual machine that is using 4000 MHz, the host is using two CPUs completely: 400 / (4 2000) = 0.50 Key: cpu|usagemhz_average
CPU|Wait	Total CPU time spent in wait state. The wait total includes time spent in the CPU Idle, CPU Swap Wait, and CPU I/O Wait states. Key:cpu|wait
CPU|Workload (%)	Percent of workload Key: cpu|workload

Memory Metrics

Memory metrics provide information about memory use and allocation.

Table 1-3. Memory Metrics

Metric Name	Description
mem|Contention (%)	This metric shows the percentage of time VMs are waiting to access swapped memory. Use this metric to monitor ESXi memory swapping. A high value indicates that the ESXi is running low on memory, and a large amount of memory is being swapped. Key: mem|host_contentionPct
mem|Machine Demand (KB)	Host memory demand in kilobytes. Key: mem|host_demand
mem|Provisioned Memory	Provisioned host memory in kilobytes. Key: mem|host_provisioned
mem|Reserved Capacity (KB)	Total amount of memory reservation used by powered-on virtual machines and vSphere services on the host. Key: mem|reservedCapacity_average
mem|Usable Memory (KB)	Usable host memory in kilobytes. Key: mem|host_usable
mem|Host Usage (KB)	Host memory use in kilobytes. Key: mem|host_usage
mem|Usage/Usable (%)	Memory usage as percentage of total configured or available memory. Key: mem|host_usagePct
mem|Workload (%)	Percent of workload. Key: mem|workload

Network Metrics

Network metrics provide information about network performance.

Table 1-4. Network Metrics

Metric Name	Description
net|Packets Dropped (%)	This metric shoes the percentage of received and transmitted packets dropped in the collection interval. Use this metric to monitor the reliability and performance of the ESXi network. A high value indicates that the network is not reliable and performance decreases. Key: net|droppedPct
net|Usage Rate (KB per second)	Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine. Key: net|usage_average
net|Workload (%)	Percent of workload. Key: net|workload

Disk Metrics

Disk metrics provide information about disk use.

Table 1-5. Disk Metrics

Metric Name	Description
disk|Commands per second	Average number of commands issued per second during the collection cycle. Key: disk|commandsAveraged_average
disk|Usage Rate (KB per second)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. Key: disk|usage_average
disk|Workload (%)	Percent of workload. Key: disk|workload

Summary Metrics

Summary metrics provide information about overall performance.

Table 1-6. Summary Metrics

Metric Name	Description
summary|Number of Running Hosts	Number of running hosts. Key: summary|number_running_hosts
summary|Number of Running VMs	This metric shows the number of running VMs at a given point in time. The data is sampled every five minutes. A large number of running VMs might be a reason for CPU or memory spikes because more resources are used in the host. The number of running VMs gives you a good indicator of how many requests the ESXi host must juggle. Powered off VMs are not included because they do not impact ESXi performance. A change in the number of running VMs can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, because all the VMs will fail if the ESXi crashes. Use this metric to look for a correlation between spikes in the running VMs and spikes in other metrics such as CPU contention, or memory contention. Key: summary|number_running_vms
summary|Total Number of Clusters	Total number of clusters. Key: summary|total_number_clusters

Table 1-6. Summary Metrics (Continued)

Metric Name	Description
summary|Total Number of Datastores	Total number of datastores. Key: summary|total_number_datastores
summary|Total Number of Hosts	Total number of hosts. Key: summary|total_number_hosts
summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
summary|Total Number of Datacenters	Total number of data centers. Key: summary|total_number_datacenters
summary|Number VCPUs on Powered on VMs	Number of virtual CPUs on powered-on virtual machines. Key: summary|number_running_vcpus
summary|Average Running VM Count per Running Host	Average running virtual machine count per running host. Key: summary|avg_vm_density

vCenter Server Metrics

vRealize Operations Manager collects CPU use, disk, memory, network, and summary metrics for vCenter Server system objects.

vCenter Server metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

CPU Usage Metrics

CPU usage metrics provide information about CPU use.

Table 1-7. CPU Usage Metrics

Metric Key	Metric Name	Description
cpu|capacity_usagepct_average	Capacity Usage (%)	Percent capacity used.
cpu|capacity_contentionPct	CPU Contention (%)	Percent CPU contention.
cpu|demandPct	Demand (%)	Percent demand.
cpu|demandmhz	Demand (MHz)	Demand in megahertz.
cpu|demand_average	Demand	CPU Demand.
cpu|iowait	IO Wait (ms)	IO wait time in milliseconds.
cpu|numpackages	Number of CPU Sockets	Number of CPU sockets.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|capacity_provisioned	Provisioned Capacity (MHz)	Provisioned capacity in megahertz.
cpu|corecount_provisioned	Provisioned vCPU	Number of provisioned virtual CPU cores.
cpu|reservedCapacity_average	Reserved Capacity (MHz)	Sum of the reservation properties of the immediate children of the host's root resource pool.
cpu|usagemhz_average	Usage (MHz)	Average CPU use in megahertz.
cpu|wait	Wait (ms)	CPU time spent on the idle state.
cpu|overhead_average	Overhead	Amount of CPU that is overhead.

Table 1-7. CPU Usage Metrics (Continued)

Metric Key	Metric Name	Description
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|vm_capacity_provisioned	Provisioned Capacity	Provisioned capacity (MHz).

Datastore Metrics

Datastore metrics provide information about the datastore.

Table 1-8. Datastore Metrics

Metric Key	Metric Name	Description
datastore| maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore| maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Maximum observed number of outstanding IO operations.
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore| numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore| numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.

Disk Metrics

Disk metrics provide information about disk use.

Table 1-9. Disk Metrics

Metric Key	Metric Name	Description
disk|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection cycle.
disk|totalLatency_average	Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Device Command Latency and Physical Device Command Latency metrics.
disk|usage_average	Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.
disk|sum_queued_oio	Total queued outstanding operations	Sum of queued operations and outstanding operations.
disk|max_observed	Max Observed OIO	Max observed IO for a disk.

Diskspace Metrics

Disk space metrics provide information about disk space use.

Table 1-10. Diskspace Metrics

Metric Key	Metric Name	Description
diskspace|total_usage	Total disk space used (KB)	Total disk space used on all datastores visible to this object.
diskspace|total_capacity	Total disk space (KB)	Total disk space on all datastores visible to this object.
diskspace|total_provisioned	Total provisioned disk space (KB)	Total provisioned disk space on all datastores visible to this object.

Memory Metrics

Memory metrics provide information about memory use and allocation.

Table 1-11. Memory Metrics

Metric Key	Metric Name	Description
mem|host_contentionPct	Contention (%)	Percent host memory contention.
mem|host_demand	Machine Demand (KB)	Host memory demand in kilobytes.
mem|host_systemUsage	ESX System Usage	Memory usage by the VMkernel and ESX user-level services.
mem|host_provisioned	Provisioned Memory (KB)	Provisioned host memory in kilobytes.
mem|reservedCapacity_average	Reserved Capacity (KB)	Sum of the reservation properties of the immediate children of the host's root resource pool.
mem|host_usable	Usable Memory (KB)	Usable host memory in kilobytes.
mem|host_usage	Host Usage (KB)	Host memory use in kilobytes.
mem|host_usagePct	Usage/Usable (%)	Percent host memory used.
mem|host_contention	Contention (KB)	Host contention in kilobytes.
mem|overhead_average	VM Overhead (KB)	Memory overhead reported by host.

Network Metrics

Network metrics provide information about network performance.

Table 1-12. Network Metrics

Metric Key	Metric Name	Description
net|droppedPct	Packets Dropped (%)	Percent network packets dropped.
net|usage_average	Usage Rate (KBps)	Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine.
net|packetsRx_summation	Packets Received	Number of packets received in the performance interval.
net|packetsTx_summation	Packets Transmitted	Number of packets transmitted in the performance interval.
net|droppedRx_summation	Received Packets Dropped	Number of received packets dropped in the performance interval.
net|droppedTx_summation	Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval.

Table 1-12. Network Metrics (Continued)

Metric Key	Metric Name	Description
net|maxObserved_KBps	Max Observed Throughput (KBps)	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput (KBps)	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput (KBps)	Max observed received rate of network throughput.
net|transmitted_average	Data Transmit Rate (KBps)	Average amount of data transmitted per second.
net|received_average	Data Receive Rate (KBps)	Average amount of data received per second.

Summary Metrics

Summary metrics provide information about overall performance.

Table 1-13. Summary Metrics

Metric Key	Metric Name	Description
summary|number_running_hosts	Number of Running Hosts	Number of hosts that are on.
summary|number_running_vms	Number of Running VMs	Number of virtual machines that are on.
summary|total_number_clusters	Total Number of Clusters	Total number of clusters.
summary|total_number_datastores	Total Number of Datastores	Total number of datastores.
summary|total_number_hosts	Total Number of Hosts	Total number of hosts.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.
summary|workload_indicator	Workload Indicator (%)	Percent workload indicator.
summary|total_number_datacenters	Total Number of Datacenters	Total number of datacenters.
summary|number_powered_on_cores	Number of Cores on Powered On Hosts	Number of cores on powered-on hosts.
summary|number_running_vcpus	Number VCPUs on Powered on VMs	Number of virtual CPUs on powered-on virtual machines.
summary|avg_vm_density	Average Running VM Count per Running Host	Average running virtual machine count per running host.
summary|vc_query_time	VC Query Time (ms)	vCenter Server query time in milliseconds.
summary|derived_metrics_comp_time	Derived Metrics Computation Time (ms)	Derived metrics computation time in milliseconds.
summary|number_objs	Number of objects	Number of objects.
summary|number_vc_events	Number of VC Events	Number of vCenter Server events.
summary|number_sms_metrics	Number of SMS Metrics	Number of SMS metrics.
summary|collector_mem_usage	Collector Memory Usage (MB)	Collector memory use in megabytes.

Virtual Machine Metrics

vRealize Operations Manager collects configuration, CPU use, memory, datastore, disk, virtual disk, guest file system, network, power, disk space, storage, and summary metrics for virtual machine objects.

Capacity metrics can be calculated for virtual machine objects. See "Capacity and Project-Based Metrics," on page 77.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the virtual machines in your environment.

Configuration Metrics for Virtual Machines

Configuration metrics provide information about virtual machine configuration.

Metric	Description
Config|Thin Provisioned Disk	Thin Provisioned Disk. Key: config|hardware|thin_Enabled
Config|Number of CPUs	Number of CPUs for a Virtual Machine. Key: config|hardware|num_Cpu
Config|Disk Space	Disk space metrics. Key: config|hardware|disk_Space

CPU Usage Metrics for Virtual Machines

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|IO Wait (ms)	CPU time spent waiting for IO. Key: cpu|iowait
CPU|Wait (ms)	Wait time in milliseconds. Key: cpu|wait
CPU|Overall CPU Contention (ms)	The amount of time the CPU cannot run due to contention. Key: cpu|capacity_contention
CPU|Reservation Used	CPU Reservation Used. Key: cpu|reservation_used
CPU|Effective Limit	CPU Effective Limit. Key: cpu|effective_limit
CPU|Estimated Entitlement	CPU Estimated Entitlement. key: cpu|estimated_entitlement
CPU|Idle (%)	Percentage time that CPU is idle. Key: cpu|idlePct
CPU|IO Wait (%)	Percentage IO Wait. Key: cpu|iowaitPct
CPU|Swap wait (%)	Percentage swap wait for CPU. Key: cpu|swapwaitPct
CPU|Wait (%)	Percentage of total CPU time spent in wait state. Key: cpu|waitPct
CPU|System (%)	Percentage CPU time spent on system processes. Key: cpu|systemSummationPct

Metric	Description
CPU|Demand Over Limit (MHz)	Amount of CPU Demand that is over the configured CPU Limit. Key: cpu|demandOverLimit
CPU|Demand Over Capacity (MHz)	Amount of CPU Demand that is over the configured CPU Capacity. Key: cpu|demandOverCapacity
CPU|Recommended Size Reduction (%)	Percentage of recommended CPU size reduction. Key: cpu|sizePctReduction
CPU|Normalized Co-stop	Percentage of co-stop time, normalized across all vCPUs. Key: cpu|perCpuCoStopPct
CPU|Recommended number of vCPUs to Add	Recommended number of vCPUs to Add to the VM. Key: cpu|numberToAdd
CPU|Recommended number of vCPUs to Remove	Recommended number of vCPUs to Remove from the VM. Key: cpu|numberToRemove
CPU|Capacity entitlement (MHz)	CPU entitlement for the VM after taking limits into account. Key: cpu|capacity_entitlement
CPU|Provisioned CPU Cores	Number of provisioned CPU cores. Key: cpu|corecount_provisioned
CPU|Capacity Demand Entitlement (%	Percent capacity demand entitlement. Key: cpu|capacity_demandEntitlementPct
* CPU|CPU Contention (%)	CPU contention as a percentage of 20-second collection interval. Key: cpu|capacity_contentionPct
CPU|Capacity Provisioned	Provisioned CPU capacity in megahertrz. Key: cpu|capacity_provisioned
CPU|Demand (MHz)	CPU demand in megahertz. Key: cpu|demandmhz
CPU|Host demand for aggregation	Host demand for aggregation. Key: cpu|host_demand_for_aggregation
CPU|Demand (ms)	The total CPU time that the VM could use if there was no contention. Key: cpu|demand_average
CPU|Demand (%)	CPU demand as a percentage of the provisioned capacity. Key: cpu|demandPct
CPU|Dynamic Entitlement	CPU Dynamic Entitlement. Key: cpu|dynamic_entitlement
* CPU|Usage (%)	This metric indicates the percentage of CPU that was used out of all the CPU that was allocated to the VM. CPU usage can indicate if the VM is undersized. Key: cpu|usage_average
CPU|Usage (MHz)	CPU use in megahertz. Key: cpu|usagemhz_average
CPU|System (ms)	CPU time spent on system processes. Key: cpu|system_summation

Metric	Description
CPU|Wait (ms)	Total time that a virtual CPU can not be run. It could be idle (halted) or waiting for an external event such as I/O. Key: cpu|wait_summation
CPU|Ready (ms)	CPU time spent in the ready state. Key: cpu|ready_summation
* CPU|Ready (%)	This metric indicates the percentage of time in which the VM was waiting in line to use the CPU on the host. A large ready time for a VM indicates that the VM needed CPU resources but the infrastructure was busy serving other VMs. This might indicate that the host is trying to serve too many VMs. Whenever the CPU ready is larger than 10%, you should check if the host is overloaded, or if the VM really needs all the resources that were allocated to it. Key: cpu|readyPct
CPU|Used (ms)	CPU time that is used. Key: cpu|used_summation
CPU|Extra (ms)	Extra CPU time in milliseconds. Key: cpu|extra_summation
CPU|Guaranteed (ms)	CPU time that is guaranteed for the virtual machine. Key: cpu|guaranteed_latest
CPU|Swap Wait (ms)	Swap wait time in milliseconds. Key: cpu|swapwait_summation
CPU|Co-stop (ms)	Time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costop_summation
CPU|Co-stop (%)	Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costopPct
CPU|Idle (ms)	CPU time that is idle. Key: cpu|idle_summation
CPU|Latency	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpu|latency_average
CPU|Max Limited	Time the VM is ready to run, but is not run due to maxing out its CPU limit setting. Key: cpu|maxlimited_summation
CPU|Overlap	Time the VM was interrupted to perform system services on behalf of that VM or other VMs. Key: cpu|overlap_summation
CPU|Run	Time the VM is scheduled to run. Key: cpu|run_summation
CPU|Entitlement Latest	Entitlement Latest. Key: cpu|entitlement_latest

CPU Utilization for Resources Metrics for Virtual Machines

CPU utilization for resources metrics provide information about resource CPU use.

Metric	Description
rescpu|CPU Active (%) (*interval*)	The average active time (actav) or peak active time (actpk) for the CPU during various intervals. Key: rescpu|actav1_latest rescpu|actav5_latest rescpu|actav15_latest rescpu|actpk1_latest rescpu|actpk5_latest rescpu|actpk15_latest
rescpu|CPU Running (%) (*interval*)	The average runtime (runav) or peak active time (runpk) for the CPU during various intervals. Key: rescpu|runav1_latest rescpu|runav5_latest rescpu|runav15_latest rescpu|runpk1_latest rescpu|runpk5_latest rescpu|runpk15_latest
rescpu|CPU Throttled (%) (*interval*)	Amount of CPU resources over the limit that were refused, average over various intervals. Key: rescpu|maxLimited1_latest rescpu|maxLimited5_latest rescpu|maxLimited15_latest
rescpu|Group CPU Sample Count	The sample CPU count. Key: rescpu|sampleCount_latest
rscpu|Group CPU Sample Period (ms)	The sample period. Key: rscpu|samplePeriod_latest

Memory Metrics for Virtual Machines

Memory metrics provide information about memory use and allocation.

Metric	Description
Mem|Host Active (KB)	Host active memory use in kilobytes. Key: mem|host_active
Mem|Usage (KB)	Memory use in kilobytes. Key: mem|host_usage
Mem|Contention (KB)	Memory contention in kilobytes. Key: mem|host_contention
Mem|Contention (%)	Percent memory contention. Key: mem|host_contentionPct
Mem|Guest Configured Memory (KB)	Guest operating system configured memory in kilobytes. Key: mem|guest_provisioned
Mem|Guest Dynamic Entitlement (KB)	Guest Memory Dynamic Entitlement. Key: mem|guest_dynamic_entitlement
Mem|Guest Active Memory (%)	Percent guest operating system active memory. Key: mem|guest_activePct

Metric	Description
Mem|Guest Non Pageable Memory (KB)	Guest operating system non-pageable memory in kilobytes. Key: mem|guest_nonpageable_estimate
Mem|Reservation Used	Memory Reservation Used. Key: mem|reservation_used
Mem|Effective Limit	Memory Effective Limit. Key: mem|effective_limit
Mem|Estimated Entitlement	Memory Estimated Entitlement. Key: mem|estimated_entitlement
Mem|Demand for aggregation	Host demand for aggregation. Key: mem|host_demand_for_aggregation
Mem|NUMA Remote Latest	Non-uniform memory access Remote (Kb). Key: mem|numa.remote_latest
Mem|NUMA Local Latest	Non-uniform memory access Local (Kb). Key: mem|numa.local_latest
Mem|NUMA Migrations Latest	Non-uniform memory access Migrations (number). Key: mem|numa.migrations_latest
Mem|NUMA Locality Average	Non-uniform memory access Locality (%). Key: mem|numa.locality_average
Mem|Demand Over Limit	Amount of Memory Demand that is over the configured Memory Limit. Key: mem|demandOverLimit
Mem|Demand Over Capacity	Amount of Memory Demand that is over the configured Memory Capacity. Key: mem|demandOverCapacity
Mem|Recommended Size Reduction (%)	Percentage of recommended Memory size reduction. Key: mem|sizePctReduction
Mem|Balloon (%)	Percentage of total memory that has been reclaimed via ballooning. Key: mem|balloonPct
* Mem|Guest Usage (KB)	This metric shows the amount of memory the VM uses. Key: mem|guest_usage
Mem|Guest Demand (KB)	Guest operating system demand in kilobytes. Key: mem|guest_demand
Mem|Guest Non Pageable Memory (KB)	Guest operating system non-pageable memory in kilobytes. Key: mem|host_nonpageable_estimate
Mem|Host Demand (KB)	Memory demand in kilobytes. Key mem|host_demand
Mem|Demand with Reservation (KB)	Memory Demand with Reservation considered in KB. Key: mem|host_demand_reservation
Mem|Guest Workload	Guest Workload (%). Key: mem|guest_workload
Mem|Host Workload	Host Workload (%). Key: host_workload

Metric	Description
Mem|Balloon (%)	Amount of memory currently used by the virtual machine memory control. Key: mem|vmmemctl_average
Mem|Guest Active (%)	Amount of memory that is actively used. Key: mem|active_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Shared (KB)	Amount of shared memory in kilobytes. Key: mem|shared_average
Mem|Zero (KB)	Amount of memory that is all 0. Key: mem|zero_average
* Mem|Swapped (KB)	This metric shows how much memory is being swapped. Meaning, the amount of unreserved memory in kilobytes. Key: mem|swapped_average
Mem|Swap Target (KB)	Amount of memory that can be swapped in kilobytes. Key: mem|swaptarget_average
Mem|Swap In (KB)	Swap-in memory in kilobytes. Key: mem|swapin_average
Mem|Swap Out (KB)	amount of memory swapped out in kilobytes. Key: mem|swapout_average
* Mem|Usage (%)	This metric shows how much memory is being used out of the allocated memory for the VM. Key: mem|usage_average
Mem|Balloon Target (KB)	Amount of memory that can be used by the virtual machine memory control. Key: mem|vmmemctltarget_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory in kilobytes. Key: mem|consumed_average
Mem|Overhead (KB)	Memory overhead in kilobytes. Key: mem|overhead_average
Mem|Host Dynamic Entitlement	Mem Machine Dynamic Entitlement. Key: mem|host_dynamic_entitlement
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem|swapinRate_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory to disk during the current interval. Key: mem|swapoutRate_average
Mem|Active Write (KB)	Active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Compressed memory in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Compression rate in kilobytes per second. Key: mem|commpressionRate_average
Mem|Decompression Rate (KBps)	Decompression rate in kilobytes per second. Key: mem|decompressionRate_average

Metric	Description
Mem|Overhead Max (KB)	Maximum overhead in kilobytes. Key: mem|overheadMax_average
Mem|Zip Saved (KB)	Zip-saved memory in kilobytes. Key: mem|zipSaved_latest
Mem|Zipped (KB)	Zipped memory in kilobytes. Key: mem|zipped_latest
Mem|Entitlement	Amount of host physical memory the VM is entitled to, as determined by the ESX schedule. Key: mem|entitlement_average
Mem|Latency	Percentage of time the VM is waiting to access swapped or compressed memory. Key: mem|latency_average
Mem|Capacity Contention	Capacity Contention. Key: mem|capacity.contention_average
Mem|Swap In Rate from Host Cache	Rate at which memory is being swapped from host cache into active memory. Key: mem|llSwapInRate_average
Mem|Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem|llSwapOutRate_average
Mem|Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem|llSwapUsed_average
Mem|Overhead Touched	Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM. Key: mem|overheadTouched_average

Datastore Metrics for Virtual Machines

Datastore metrics provide information about datastore use.

Metric	Description
Datastore|Commands per second	Average number of commands issued per second during the collection interval. Key:datastore|commandsAveraged_average
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore|oio
Datastore|Demand	Datastore demand. Key: datastore|demand
Datastore|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Usage Average (KBps). Key: datastore|usage_average
Datastore|Used Space (MB)	Used space in megabytes. Key: datastore|used

Metric	Description
Datastore|Not Shared (GB)	Space used by VMs that is not shared. Key: datastore|notshared
* Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
* Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average
* Datastore|Read Rate (KBps)	This metric shows the amount of data that the VM reads to the datastore per second.. Key: datastore|read_average
*Datastore|Read Latency (ms)	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
*Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average
* Datastore|Write Rate	This metric shows the amount of data that the VM writes to the datastore per second. Key: datastore|write_average
Datastore|Highest Latency	Highest Latency. Key: datastore|maxTotalLatency_latest
Datastore|Total Latency Max	Total Latency Max (ms). Key: datastore|totalLatency_max
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO

Disk Metrics for Virtual Machines

Disk metrics provide information about disk use.

Metric	Description
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_average
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average

Metric	Description
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Usage Rate (KBps)	Use rate in kilobytes per second. Key: disk|usage_average
Disk|I/O Usage Capacity	This metric is a function of storage|usage_average and disk|workload. storage|usage_average is an average over all storage devices. This means that disk|usage_capacity is not specific to the selected VM or the host of the VM. Key: disk|usage_capacity
Disk|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: disk|diskoio
Disk|Queued Operations	Queued operations. Key: disk|diskqueued
Disk|Demand (%)	Percent demand. Key: disk|diskdemand
Disk |Total Queued Outstanding Operations	Sum of Queued Operation and Outstanding Operations. Key: disk |sum_queued_oio
Disk|Max Observed OIO	Max Observed IO for a disk. Key: disk|max_observed
Disk|Read Rate (KBps)	Amount of data read in the performance interval. Key: disk|read_average
Disk|Write Rate (KBps)	Amount of data written to disk in the performance interval. Key: disk|write_average
Disk|Read Requests	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Write Requests	Number of times data was written to the disk in the defined interval. Key: disk|numberWrite_summation
Disk|Bus Resets	The number of bus resets in the performance interval. Key: disk|busResets_summation
Disk|Commands Issued	The number of disk commands issued in the performance interval. Key: disk|commands_summation
Disk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: disk|commandsAborted_summation
Disk|Highest Latency	Highest latency. Key: disk|maxTotalLatency_latest
Disk|SCSI Reservation Conflicts	SCSI Reservation Conflicts. Key: disk|scsiReservationConflicts_summation
Disk|Disk Read Latency	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalReadLatency_average

Metric	Description
Disk|Disk Write Latency	The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency. Key: disk|totalWriteLatency_average
Disk|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: disk|totalLatency_average

Virtual Disk Metrics for Virtual Machines

Virtual disk metrics provide information about virtual disk use.

Metric	Description
VirtualDisk|Usage	Average CPU usage as a percentage. Key: virtualDisk|usage
VirtualDisk|Total Latency	Total latency. Key: virtualDisk|totalLatency
VirtualDisk|Commands Per Second	Average number of commands per second. Key: virtualDisk|commandsAveraged_average
VirtualDisk|Read Requests	Average number of read commands issued per second to the virtual disk during the collection interval. Key: virtualDisk|numberReadAveraged_average
VirtualDisk|Write Requests	Average number of write commands issued per second to the virtual disk during the collection interval. Key: virtualDisk|numberWriteAveraged_average
VirtualDisk|Read Rate (KBps)	Rate of reading data from the virtual disk in kilobytes per second. Key: virtualDisk|read_average
VirtualDisk|Read Latency (ms)	Average amount of time for a read operation from the virtual disk. Total latency = kernel latency + device latency. Key: virtualDisk|totalReadLatency_average
VirtualDisk|Write Latency (ms)	Average amount of time for a write operation to the virtual disk. Total latency = kernel latency + device latency. Key: virtualDisk|totalWriteLatency_average
VirtualDisk|Write Rate (KBps)	Rate of writing data from the virtual disk in kilobytes per second. Key: virtualDisk|write_average
VirtualDisk|Bus Resets	The number of bus resets in the performance interval. Key: virtualDisk|busResets_summation
VirtualDisk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: virtualDisk|commandsAborted_summation
VirtualDisk|Read Load	Storage DRS virtual disk metric read load. Key: virtualDisk|readLoadMetric_latest
VirtualDisk|Outstanding Read Requests	Average number of outstanding read requests to the virtual disk. Key: virtualDisk|readOIO_latest

Metric	Description
VirtualDisk|Write Load	Storage DRS virtual disk write load. Key: virtualDisk|writeLoadMetric_latest
VirtualDisk|Outstanding Write Requests	Average number of outstanding write requests to the virtual disk. Key: virtualDisk|writeOIO_latest
VirtualDisk|Number of Small Seeks	Small Seeks. Key: virtualDisk|smallSeeks_latest
VirtualDisk|Number of Medium Seeks	Medium Seeks. Key: virtualDisk|mediumSeeks_latest
VirtualDisk|Number of Large Seeks	Large Seeks. Key: virtualDisk|largeSeeks_latest
VirtualDisk|Read Latency (microseconds)	Read Latency in microseconds. Key: virtualDisk|readLatencyUS_latest
VirtualDisk|Write Latency (microseconds)	Write Latency in microseconds. Key: virtualDisk|writeLatencyUS_latest
VirtualDisk|Average Read request size	Read IO size. Key: virtualDisk|readIOSize_latest
VirtualDisk|Average Write request size	Write IO size. Key: virtualDisk|writeIOSize_latest

Guest File System Metrics for Virtual Machines

Guest file system metrics provide information about guest file system capacity and free space.

Metric	Description
Guest file system|Guest File System Capacity (MB)	Total capacity on guest file system in megabytes. Key: guestfilesystem|capacity
Guest file system|Guest File System Free (MB)	Total free space on guest file system in megabytes. Key: guestfilesystem|freespace
Guest file system|Guest File System Usage (%)	Percent guest file system. Key: guestfilesystem|percentage
Guest file system|Guest File System Usage	Total usage of guest file system. Key: guestfilesystem|usage
* Guest file system|Total Guest File System Free (GB)	This metric displays the amount of free disk space from all file systems attached to this VM. Use this metric to see if there are spikes in the free space or if there is an organic growth for this VM. Key: guestfilesystem|freespace_total
* Guest file system|Total Guest File System Capacity (GB)	This metric displays the amount of disk space allocated for the VM. Correlate other metrics with this metric to indicate if changes occur in the disk space allocation for the VM. Key: guestfilesystem|capacity_total

Metric	Description
* Guest file system|Total Guest File System Usage (%)	This metric displays the amount of displa space being used out of the total allocated disk space. Use his metric to track if the overall usage is stable, or if it reaches the limits. You should avoid having VMs with a disk space usage of >95% since this could impact your system. Key: guestfilesystem|percentage_total
Guest file system|Total Guest File System Usage	Total usage of guest file system. Key: guestfilesystem|usage_total

Network Metrics for Virtual Machines

Network metrics provide information about network performance.

Metric	Description
Net|Demand (%)	Percent demand. Key: net|demand
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|Packets Received per second	Number of packets received in the performance interval. Key: net|packetsRxPerSec
Net|Packets Transmitted per second	Number of packets transmitted in the performance interval. Key: net|packetsTxPerSec
* Net|Data Transmit Rate (KBps)	This metric shows the rate of data being sent by the VM per second. Key: net|transmitted_average
* Net|Data Receive Rate (KBps)	This metric shows the rate of data received by the VM per second. Key: net|received_average
Net|Packets per second	Number of packets transmitted and received per second. Key: net|PacketsPerSec
Net|I/O Usage Capacity	IO use capacity. Key: net|usage_capacity
Net|Max Observed Throughput (KBps)	Maximum observed throughput in kilobytes per second. Key: net|maxObserved_KBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
* Net|Received Packets Dropped	This metric shows the number of received packets dropped in the collection interval. Key: net|droppedRx_summation
* Net|Transmitted Packets Dropped	This metric shows the number of transmitted packets dropped in the collection interval Key: net|dropppedTx_summation

Metric	Description
Net|Packets Dropped (%)	Percentage of packets dropped. Key: net|droppedPct
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|Broadcast Packets Transmitted	Number of broadcast packets transmitted during the sampling interval. Key: net|broadcastTx_summation
Net|Broadcast Packets Received	Number of broadcast packets received during the sampling interval. Key: net|broadcastRx_summation
Net|Bytes Rx (KBps)	Average amount of data received per second. Key: net|bytesRx_average
Net|Bytes Tx (KBps)	Average amount of data transmitted per second. Key: net|bytesTx_average
Net|Multicast Packets Received	Number of multicast packets received. Key: net|multicastRx_summation
Net|Multicast Packets Transmitted	Number of multicast packets transmitted. Key: net|multicastTx_summation
Net|VM to Host Data Transmit Rate	Average amount of data transmitted per second between VM and host. Key: net|host_transmitted_average
Net|VM to Host Data Receive Rate	Average amount of data received per second between VM and host. Key: net|host_received_average
Net|VM to Host Usage Rate	The sum of the data transmitted and received for all the NIC instances between VM and host. Key: net|host_usage_average
Net|VM to Host Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput between VM and host. Key: net|host_maxObserved_Tx_KBps
Net|VM to Host Max Observed Received Throughput	Max observed received rate of network throughput between VM and host. Key: net|host_maxObserved_Rx_KBps
Net|VM to Host Max Observed Throughput	Max observed rate of network throughput between VM and host. Key: net|host_maxObserved_KBps
Net|Data Transmit Demand Rate	Data transmitted demand rate. Key: net|transmit_demand_average
Net|Data Receive Demand Rate	Data received demand rate. Key: net|receive_demand_average

System Metrics for Virtual Machines

System metrics for virtual machines provide general information about the virtual machine, such as its build number and running state.

Metric	Description
Sys|Powered ON	Powered on virtual machines. 1 if powered on, 0 if powered off, -1 if unknown Key: sys|poweredOn
Sys|Uptime (seconds)	Number of seconds since system startup. Key: sys|uptime_latest
Sys|Heartbeat	Number of heartbeats from the virtual machine in the defined interval. Key: sys|heartbeat_summation
Sys|vMotion Enabled	1 if vMotion is enabled or 0 if vMotion is not enabled. Key: sys|vmotionEnabled
Sys|Product String	VMware product string. Key: sys|productString
Sys|Build Number	VMware build number. Key: sys|build
Sys|OS Uptime	Total time elapsed, in seconds, since last operating system boot-up. Key: sys|osUptime_latest

Power Metrics for Virtual Machines

Power metrics provide information about power use.

Metric	Description
Power|Energy (Joule)	Energy use in joules. Key: power|energy_summation
Power|Power (Watt)	Average power use in watts. Key: power|power_average

Disk Space Metrics for Virtual Machines

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared space in kilobytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disks	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Provisioned Space (GB)	Provisioned space in gigabytes. Key: diskspace|provisioned
Diskspace|Provisioned Space for VM	Provisioned space for VM. Key: diskspace|provisionedSpace
Diskspace|Shared Used (GB)	Shared used space in gigabytes. Key: diskspace|shared
Diskspace|Snapshot Space (GB)	Space used by snapshots. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Space used by virtual disks in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Space used by virtual machine files in gigabytes. Key: diskspace|used

Metric	Description
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object.Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned
Diskspace|Active not shared	Unshared disk space used by VMs excluding snapshot. Key: diskspace|activeNotShared

Storage Metrics for Virtual Machines

Storage metrics provide information about storage use.

Metric	Description
Storage|Commands per second	Average number of commands issued per second during the collection interval. Key: storage|commandsAveraged_average
Storage|Contention (%)	Percent contention. Key: storage|contention
Storage|Demand (KBps)	Demand in kilobytes per second. Key: storage|demandKBps
* Storage|Read Latency (ms)	This metric shows the latency that the VM experiences while performing a read action. Key: storage|totalReadLatency_average
Storage|Read Rate (KBps)	Read throughput rate in kilobytes per second. Key: storage|read_average
Storage|Reads per second	Average number of read commands issued per second during the collection interval. Key: storage|numberReadAveraged_average
Storage|Total Latency (ms)	Total latency in milliseconds. Key: storage|totalLatency_average
Storage|Total Usage (KBps)	Total throughput rate in kilobytes per second. Key: storage|usage_average
* Storage|Write Latency (ms)	This metric shows the latency that this VM experiences while performing a write action. Key: storage|totalWriteLatency_average
Storage|Write Rate (KBps)	Write throughput rate in kilobytes per second. Key: storage|write_average
Storage|Writes per second	Average number of write commands issued per second during the collection interval. Key: storage|numberWriteAveraged_average

Summary Metrics for Virtual Machines

Summary metrics provide information about overall performance.

Metric	Description
Summary|Workload Indicator (%)	Percent workload indicator. Key: summary|workload_indicator
Summary|CPU Shares	CPU shares. Key: summary|cpu_shares
Summary|Memory Shares	Memory shares. Key: summary|mem_shares
Summary|Number of Datastores	Number of datastores. Key: summary|number_datastore
Summary|Number of Networks	Number of networks. Key: summary|number_network
Summary|Running	Number of running virtual machines. Key: summary|running
Summary|Desktop Status	Horizon View Desktop Status. Key: summary|desktop_status

Host System Metrics

vRealize Operations Manager collects many metrics for host systems, including CPU use, datastore, disk, memory, network, storage, and summary metrics for host system objects.

Capacity metrics can be calculated for host system objects. See "Capacity and Project-Based Metrics," on page 77.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the hosts in your environment.

vFlash Module Metrics for Host Systems

vFlash Module metrics provide information about the host system's flash devices.

Metric	Description
vFlashModule|Latest Number of Active VM Disks	Latest Number of Active VM Disks. Key: vflashModule|numActiveVMDKs_latest

Configuration Metrics for Host Systems

Configuration metrics provide information about host system configuration.

Metric	Description
Configuration|Failover Hosts	Failover Hosts. Key: configuration|dasConfig|admissionControlPolicy|failoverHost

Hardware Metrics for Host Systems

Hardware metrics provide information about host system hardware.

Metric	Description
Hardware|Number of CPUs	Number of CPUs for a host. Key: hardware|cpuinfo|num_CpuCores

CPU Usage Metrics for Host Systems

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|Capacity Usage (%)	Percent CPU capacity used. Key: cpu|capacity_usagepct_average
CPU|Usage (%)	Average CPU usage as a percentage. Key: cpu|usage_average
* CPU|CPU Contention (%)	This metric indicates the percentage of time the virtual machines in the ESXi hosts are unable to run because they are contending for access to the physical CPU(s). This is the average number of all VMs. Naturally, the number will be lower than the highest number experienced by the worst hit VM (a VM that suffers the highest CPU contention). Use this metric to verify if the host is able to serve all of its VMs well. When using this metric, ensure the number is within your expectation. The metric is affected by several factors so you need to watch both relative numbers and absolute numbers. Relative means a drastic change in value. This indicates that the ESXi is unable to service its VMs. Absolute means that the real value is high. You should investigate why the value is high. One factor that impacts the CPU contention metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz that reduction in speed is taken into account. This is because the VM is not running at full speed. Key: cpu|capacity_contentionPct
* CPU|Demand (%)	This metric shows the percentage of CPU resources all the VMs would use if there was no CPU contention or any CPU limits set. It represents the average active CPU load for the past five minutes. Keep the number of this metric below 100% if you set Power Management to Maximum. Key: cpu|demandPct
CPU|Demand (MHz)	CPU demand in megahertz. Key: cpu|demandmhz
CPU|IO Wait (ms)	IO wait time in milliseconds. Key: cpu|iowait
CPU|Number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention (ms)	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Provisioned Capacity (MHz)	Capacity in MHz of the physical CPU cores. Key: cpu|capacity_provisioned
CPU|Provisioned virtual CPUs	Provisioned virtual CPUs. Key: cpu|corecount_provisioned
CPU|Total Wait	CPU time spent in idle state. Key: cpu|wait
CPU|Demand	CPU demand. Key: cpu|demand_average

Metric	Description
CPU|Used (msec)	Time accounted to the virtual machine. If a system service runs on behalf of this virtual machine, the time spent by that service (represented by cpu.system) should be charged to this virtual machine. If not, the time spent (represented by cpu.overlap) should not be charged against this virtual machine. Key: cpu|used_summation
CPU|Usage (MHz)	CPU use in megahertz. Key: cpu|usagemhz_average
CPU|Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool. Key: cpu|reservedCapacity_average
CPU|Total Capacity (MHz)	Total CPU capacity in megahertz. Key: cpu|totalCapacity_average
CPU|Idle (ms)	CPU idle time in milliseconds. Key: cpu|idle_summation
CPU|Overhead (KB)	Amount of CPU overhead. Key: cpu|overhead_average
CPU|Demand without overhead	Value of demand excluding any overhead. Key: cpu|demand_without_overhead
CPU|Core Utilization (%)	Percent core utilization. Key: cpu|coreUtilization_average
CPU|Utilization(%)	Percent CPU utilization. Key: cpu|utilization_average
CPU|Core Utilization (%)	Core Utilization. Key: cpu|coreUtilization_average
CPU|Utilization (%)	Utilization. Key: cpu|utilization_average
CPU|Co-stop (ms)	Time the VM is ready to run, but is unable to due to co-scheduling constraints. Key: cpu|costop_summation
CPU|Latency (%)	Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. Key: cpu|latency_average
CPU|Ready (ms)	Time spent in ready state. Key: cpu|ready_summation
CPU|Run (ms)	Time the virtual machine is scheduled to run. Key: cpu|run_summation
CPU|Swap wait (ms)	Amount of time waiting for swap space. Key: cpu|swapwait_summation
CPU|Wait (ms)	Total CPU time spent in wait state. Key: cpu|wait_summation
CPU|Provisioned Capacity	Provisioned capacity (MHz). Key: cpu|vm_capacity_provisioned

Metric	Description
CPU|Active Host Load For Balance (Long Term)	Active Host Load For Balance (Long Term). Key: cpu|acvmWorkloadDisparityPcttive_longterm_load
CPU|Active Host Load For Balance (Short Term)	Active Host Load For Balance (Short Term). Key: cpu|active_shortterm_load

CPU Utilization for Resources Metrics for Host Systems

CPU utilization for resources metrics provide information about CPU activity.

Metric Name	Description
Rescpu|CPU Active (%) (*interval*)	Average active time for the CPU over the past minute, past five minutes, and at one-minute, five-minute, and 15-minute peak active times. Key: rescpu|actav1_latest rescpu|actav5_latest rescpu|actav15_latest rescpu|actpk1_latest rescpu|actpk5_latest rescpu|actpk15_latest
Rescpu|CPU Running (%) (*interval*)	Average run time for the CPU over the past minute, past five minutes, past 15 minutes, and at one-minute, five-minute, and 15-minute peak times. Key: rescpu|runav1_latest rescpu|runav5_latest rescpu|runav15_latest rescpu|runpk1_latest rescpu|runpk5_latest rescpu|runpk15_latest
Rescpu|CPU Throttled (%) (*interval*)	Scheduling limit over the past minute, past five minutes, and past 15 minutes. Key: rescpu|maxLimited1_latest rescpu|maxLimited5_latest rescpu|maxLimited15_latest
Rescpu|Group CPU Sample Count	Group CPU sample count. Key: rescpu|sampleCount_latest
Rescpu|Group CPU Sample Period (ms)	Group CPU sample period in milliseconds. Key: rescpu|samplePeriod_latest

Datastore Metrics for Host Systems

Datastore metrics provide information about datastore use.

Metric	Notes
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead

Metric	Notes
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO
Datastore|Commands Averaged	Average number of commands issued per second during the collection interval. Key: datastore|commandsAveraged_average
Datastore|Number of Outstanding IO Operations	Number of outstanding IO operations. Key: datastore|oio
Datastore|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Usage Average (KBps). Key: datastore|usage_average
Datastore|Demand	Demand. Key: datastore|demand
Datastore|Storage I/O Control aggregated IOPS	Aggregate number of IO operations on the datastore. Key: datastore|datastoreIops_average
Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average
Datastore|Read Rate (KBps)	Rate of reading data from the datastore in kilobytes per second. Key: datastore|read_average
Datastore|Storage I/O Control normalized latency (ms)	Normalized latency in microseconds on the datastore. Data for all virtual machines is combined. Key: datastore| sizeNormalizedDatastoreLatency_average
Datastore|Read Latency (ms)	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average

Metric	Notes
Datastore|Write Rate (KBps)	Rate of writing data to the datastore in kilobytes per second. Key: datastore|write_average
Datastore|Max Queue Depth	Max Queue Depth. Key: datastore|datastoreMaxQueueDepth_latest
Datastore|Highest Latency	Highest Latency. Key: datastore|maxTotalLatency_latest
Datastore|Total Latency Max	Total Latency Max (ms). Key: datastore|totalLatency_max
Datastore|Read Latency	Read Latency. Key: datastore|datastoreNormalReadLatency_latest
Datastore|Write Latency	Write Latency. Key: datastore|datastoreNormalWriteLatency_latest
Datastore|Data Read	Data Read. Key: datastore|datastoreReadBytes_latest
Datastore|Data Read Rate	Data Rate. Key: datastore|datastoreReadIops_latest
Datastore|Read Load	Storage DRS metric read load. Key: datastore|datastoreReadLoadMetric_latest
Datastore|Outstanding Read Requests	Outstanding Read Requests. Key: datastore|datastoreReadOIO_latest
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|Average Observed Virtual Machine Disk I/O Workload	Average Observed Virtual Machine Disk I/O Workload on the Host. Key: datastore|vmPopulationAvgWorkload
Datastore|Maximum Observed VM Disk I/O Workload	Maximum Observed VM Disk I/O Workload on the Host. Key: datastore|vmPopulationMaxWorkload
Datastore|VM Disk I/O Workload Disparity	Percentage Disk I/O workload disparity among the VMs on the Host. Key: datastore|vmWorkloadDisparityPc

Disk Metrics for Host Systems

Disk metrics provide information about disk use.

Metric	Description
Disk|Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. disk|usage_average
Disk|I/O Usage Capacity	This metric is a function of storage|usage_average and disk|workload. storage|usage_average is an average over all storage devices. This means that disk|usage_capacity is not specific to the selected VM or the host of the VM. Key: disk|usage_capacity
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Disk Command Latency (ms)	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency. Key: disk|totalLatency_average
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_average
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average
Disk|Read Requests	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Write Requests	Number of times data was written to the disk in the defined interval. Key: disk|numberWrite_summation
Disk|Read Rate	Amount of data read in the performance interval. Key: disk|read_average
Disk|Write Rate	Amount of data written to disk in the performance interval. Key: disk|write_average
Disk|Bus Resets	The number of bus resets in the performance interval. Key: disk|busResets_summation
Disk|Commands Issued	The number of disk commands issued in the performance interval. Key: disk|commands_summation
Disk|Commands Aborted	The number of disk commands aborted in the performance interval. Key: disk|commandsAborted_summation
Disk|Physical Device Read Latency (ms)	The average time taken to complete a read from the physical device. Key: disk|deviceReadLatency_average
Disk|Kernel Disk Read Latency (ms)	The average time spent in ESX Server VMKernel per read. Key: disk|kernelReadLatency_average

Metric	Description
Disk|Disk Read Latency (ms)	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalReadLatency_average
Disk|Queue Read Latency (ms)	The average time spent in the ESX Server VMKernel queue per read. Key: disk|queueReadLatency_average
Disk|Physical Device Write Latency (ms)	The average time taken to complete a write from the physical device. Key: disk|deviceWriteLatency_average
Disk|Kernel Disk Write Latency (ms)	The average time spent in ESX Server VMKernel per write. Key: disk|kernelWriteLatency_average
Disk|Disk Write Latency (ms)	The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency. Key: disk|totalWriteLatency_average
Disk|Queue Write Latency (ms)	The average time spent in the ESX Server VMKernel queue per write. Key: disk|queueWriteLatency_average
Disk|Physical Device Command Latency (ms)	The average time taken to complete a command from the physical device. Key: disk|deviceLatency_average
Disk|Kernel Disk Command Latency (ms)	The average time spent in ESX Server VMKernel per command. Key: disk|kernelLatency_average
Disk|Queue Command Latency (ms)	The average time spent in the ESX Server VMKernel queue per command. Key: disk|queueLatency_average
Disk|Number of Outstanding IO Operations	Number of Outstanding IO Operations. Key: disk|diskoio
Disk|Queued Operations	Queued Operations. Key: disk|diskqueued
Disk|Demand	Demand. Key: disk|diskdemand
Disk|Total Queued Outstanding operations	Sum of Queued Operation and Outstanding Operations. Key: disk|sum_queued_oio
Disk|Max Observed OIO	Max Observed IO for a disk. Key: disk|max_observed
Disk|Highest Latency	Highest Latency. Key: disk|maxTotalLatency_latest
Disk|Max Queue Depth	Maximum queue depth during the collection interval. Key: disk|maxQueueDepth_average
Disk|SCSI Reservation Conflicts	SCSI Reservation Conflicts. Key: disk|scsiReservationConflicts_summation

Memory Metrics for Host Systems

Memory metrics provide information about memory use and allocation.

Metric	Description
* Mem|Contention (%)	This metric is used to monitor ESXi memory usage. When the value is high, it means the ESXi is using a good percentage of available memory. You may need to add more memory to other memory-related metrics. Key: mem|host_contentionPct
Mem|Contention (KB)	Host contention in kilobytes. Key: mem|host_contention
Mem|Host Usage (KB)	Machine usage in kilobytes. Key: mem|host_usage
Mem|Machine Demand (KB)	Host demand in kilobytes. Key: mem|host_demand
Mem|Overall Memory used to run VMs on Host (KB)	Overall memory used to run virtual machines on the host in kilobytes. Key: mem|host_usageVM
Mem|Provisioned Memory (KB)	Provisioned memory in kilobytes. Key: mem|host_provisioned
Mem|Minimum Free Memory (KB)	Minimum free memory. Key: mem|host_minfree
Mem|Reserved Capacity (%)	Percent reserved capacity. Key: mem|reservedCapacityPct
Mem|Usable Memory (KB)	Usable memory in kilobytes. Key: mem|host_usable
* Mem|Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem|host_usagePct
Mem|ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem|host_systemUsage
Mem|Guest Active (KB)	Amount of memory that is actively used. Key: mem|active_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory. Key: mem|consumed_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Heap (KB)	Amount of memory allocated for heap. Key: mem|heap_average
Mem|Heap Free (KB)	Amount of free space in the heap. Key: mem|heapfree_average
Mem|VM Overhead (KB)	Memory overhead reported by host. Key: mem|overhead_average
Mem|Reserved Capacity (KB)	Reserved capacity in kilobytes. Key: mem|reservedCapacity_average

Metric	Description
Mem|Shared (KB)	Amount of shared memory in kilobytes. Key: mem|shared_average
Mem|Shared Common (KB)	Amount of shared common memory in kilobytes. Key: mem|sharedcommon_average
Mem|Swap In (KB)	Amount of memory swapped in. Key: mem|swapin_average
Mem|Swap Out KB)	Amount of memory swapped out. Key: mem|swapout_average
Mem|Swap Used (KB)	Amount of memory used for swapped space in kilobytes. Key: mem|swapused_average
Mem|VM kernel Usage (KB)	Amount of memory used by the VM kernel. Key: mem|sysUsage_average
Mem|Unreserved (KB)	Amount of unreserved memory in kilobytes. Key: mem|unreserved_average
* Mem|Balloon (KB)	This metric shows the total amount of memory currently used by the VM memory control. This memory was reclaimed from the respective VMs at some point in the past, and was not returned. Use this metric to monitor how much VM memory has been reclaimed by ESXi through memory ballooning. The presence of ballooning indicates the ESXi has been under memory pressure. The ESXi activates ballooning when consumed memory reaches a certain threshold. Look for increasing size of ballooning. This indicates that there has been a shortage of memory more than once. Look for size fluctuations which indicate the ballooned out page was actually required by the VM. This translates into a memory performance problem for the VM requesting the page, since the page must first be brought back from the disk. Key: mem|vmmemctl_average
Mem|Zero (KB)	Amount of memory that is all zero. Key: mem|zero_average
Mem|State (0-3)	Overall state of the memory. The value is an integer between 0 (high) and 3 (low). Key: mem|state_latest
Mem|Usage (KB)	Host memory use in kilobytes. Key: mem|host_usage
Mem|Usage (%)	Memory currently in use as a percentage of total available memory. Key: mem|usage_average
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval in kilobyte per second. Key: mem|swapinRate_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory to disk during the current interval in kilobytes per second. Key: mem|swapoutRate_average

Metric	Description
Mem|Active Write (KB)	Average active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Average memory compression in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Average compression rate in kilobytes per second. Key: mem|compressionRate_average
Mem|Decompression Rate (KBps)	Decompression rate in kilobytes per second. Key: mem|decompressionRate_average
Mem|Total Capacity (KB)	Total capacity in kilobytes. Key: mem|totalCapacity_average
Mem|Latency	Percentage of time the VM is waiting to access swapped or compressed memory. Key: mem|latency_average
Mem|Capacity Contention	Capacity Contention. Key: mem|capacity.contention_average
Mem|Swap In Rate from Host Cache	Rate at which memory is being swapped from host cache into active memory. Key: mem|llSwapInRate_average
Mem|Swap In from Host Cache	Amount of memory swapped-in from host cache. Key: mem|llSwapIn_average
Mem|Swap Out Rate to Host Cache	Rate at which memory is being swapped to host cache from active memory. Key: mem|llSwapOutRate_average
Mem|Swap Out to Host Cache	Amount of memory swapped-out to host cache. Key: mem|llSwapOut_average
Mem|Swap Space Used in Host Cache	Space used for caching swapped pages in the host cache. Key: mem|llSwapUsed_average
Mem|Low Free Threshold	Threshold of free host physical memory below which ESX will begin reclaiming memory from VMs through ballooning and swapping. Key: mem|lowfreethreshold_average
Mem|VM Memory Workload Disparity	Percentage Memory workload disparity among the VMs on the Host. Key: mem|vmWorkloadDisparityPct
Mem|Active Host Load For Balance (Long Term)	Active Host Load For Balance (Long Term). Key: mem|active_longterm_load
Mem|Active Host Load For Balance (Short Term)	Active Host Load For Balance (Short Term). Key: mem|active_shortterm_load

Network Metrics for Host Systems

Network metrics provide information about network performance.

Metric	Description
* Net|Packets Received per second	This metric shows the number of packets received during the collection interval. Use this metric to monitor the network usage of the ESXi. Key: net|packetsRxPerSec
* Net|Packets Transmitted per second	This metric shows the number of packets transmitted during the collection interval. Key: net|packetsTxPerSec
Net|Packets per second	Number of packets transmitted and received per second. Key: net|packetsPerSec
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|I/O Usage Capacity	I/O Usage Capacity. Key: net|usage_capacity
Net|Max Observed Throughput	Max observed rate of network throughput. Key: net|maxObserved_KBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps
Net|Demand (%)	Percent demand. Key: net|demand
Net|Data Transmit Rate (KBps)	Average amount of data transmitted per second. Key: net|transmitted_average
Net|Data Receive Rate (KBps)	Average amount of data received per second. Key: net|received_average
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
Net|Received Packets Dropped	Number of received packets dropped in the performance interval. Key: net|droppedRx_summation
Net|Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval. Key: net|droppedTx_summation
* Net|Packets Dropped (%)	This metric shows the percentage of received and transmitted packets dropped during the collection interval. This metric is used to monitor reliability and performance of the ESXi network. When a high value is displayed, the network is not reliable and performance suffers. Key: net|droppedPct
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|bytes Rx (KBps)	Average amount of data received per second. Key: net|bytesRx_average

Metric	Description
Net|bytes Tx (KBps)	Average amount of data transmitted per second. Key: net|bytesTx_average
Net|Broadcast Packets Received	Number of broadcast packets received during the sampling interval. Key: net|broadcastRx_summation
Net|Broadcast Packets Transmitted	Number of broadcast packets transmitted during the sampling interval. Key: net|broadcastTx_summation
Net|Error Packets Received	Number of packets with errors received. Key: net|errorsRx_summation
Net|Error Packets Transmitted	Number of packets with errors transmitted. Key: net|errorsTx_summation
Net|Multicast Packets Received	Number of multicast packets received. Key: net|multicastRx_summation
Net|Multicast Packets Transmitted	Number of multicast packets transmitted. Key: net|multicastTx_summation
Net|FT Throughput Usage	FT Throughput Usage. Key: net|throughput.usage.ft_average
Net|HBR Throughput Usage	HBR Throughput Usage. Key: net|throughput.usage.hbr_average
Net|iSCSI Throughput Usage	iSCSI Throughput Usage. Key: net|throughput.usage.iscsi_average
Net|NFS Throughput Usage	NFS Throughput Usage. Key: net|throughput.usage.nfs_average
Net|VM Throughput Usage	VM Throughput Usage. Key: net|throughput.usage.vm_average
Net|vMotion Throughput Usage	vMotion Throughput Usage. Key: net|throughput.usage.vmotion_average
Net|Unknown Protocol Frames Received	Number of frames with unknown protocol received. Key: net|unknownProtos_summation

System Metrics for Host Systems

System metrics provide information about the amount of CPU that resources and other applications use.

Metric	Description
Sys|Power On	1 if the host system is powered on, 0 if the host system is powered off, or -1 if the power state is unknown. Key: sys|poweredOn
Sys|Uptime (seconds)	Number of seconds since the last system startup. Key: sys|uptime_latest
Sys|Disk Usage (%)	Percent disk use. Key: sys|diskUsage_latest
Sys|Resource CPU Usage (MHz)	Amount of CPU that the Service Console and other applications use. Key: sys|resourceCpuUsage_average

Metric	Description
Sys|Resource CPU Active (1 min. average)	Percentage of resource CPU that is active. Average value during a one-minute period. Key: sys|resourceCpuAct1_latest
Sys|Resource CPU Active (%) (5 min. average)	Percentage of resource CPU that is active. Average value during a five-minute period. Key: sys|resourceCpuAct5_latest
Sys|Resource CPU Alloc Max (MHz)	Maximum resource CPU allocation in megahertz. Key: sys|resourceCpuAllocMax_latest
Sys|Resource CPU Alloc Min (MHz)	Minimum resource CPU allocation in megahertz. Key: sys|resourceCpuAllocMin_latest
Sys|Resource CPU Alloc Shares	Number of resource CPU allocation shares. Key: sys|resourceCpuAllocShares_latest
Sys|Resource CPU Max Limited (%) (1 min. average)	Percent of resource CPU that is limited to the maximum amount. Average value during a one-minute period. Key: sys|resourceCpuMaxLimited1_latest
Sys|Resource CPU Max Limited (%) (5 min. average)	Percentage of resource CPU that is limited to the maximum amount. Average value during a five-minute period. Key: sys|resourceCpuMaxLimited5_latest
Sys|Resource CPU Run1 (%)	Percent resource CPU for Run1. Key: sys|resourceCpuRun1_latest
Sys|Resource CPU Run5 (%)	Percent resource CPU for Run5. Key: sys|resourceCpuRun5_latest
Sys|Resource Memory Alloc Max (KB)	Maximum resource memory allocation in kilobytes. Key: sys|resourceMemAllocMax_latest
Sys|Resource Memory Alloc Min (KB)	Minimum resource memory allocation in kilobytes. Key: sys|resourceMemAllocMin_latest
Sys|Resource Memory Alloc Shares	Number of resource memory shares allocated. Key: sys|resourceMemAllocShares_latest
Sys|Resource Memory Cow (KB)	Cow resource memory in kilobytes. Key: Sys|resourceMemCow_latest
Sys|Resource Memory Mapped (KB)	Mapped resource memory in kilobytes. Key: ys|resourceMemMapped_latest
Sys|Resource Memory Overhead (KB)	Resource memory overhead in kilobytes. Key: sys|resourceMemOverhead_latest
Sys|Resource Memory Shared (KB)	Shared resource memory in kilobytes. Key: sys|resourceMemShared_latest
Sys|Resource Memory Swapped (KB)	Swapped resource memory in kilobytes. Key: sys|resourceMemSwapped_latest
Sys|Resource Memory Touched (KB)	Touched resource memory in kilobytes. Key: sys|resourceMemTouched_latest
Sys|Resource Memory Zero (KB)	Zero resource memory in kilobytes. Key: sys|resourceMemZero_latest
Sys|Resource Memory Consumed	Resource Memory Consumed Latest (KB). Key: sys|resourceMemConsumed_latest
Sys|Resource File descriptors usage	Resource File descriptors usage (KB). Key: sys|resourceFdUsage_latest

Metric	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	Description
Management Agent|Memory Used (%)	Amount of total configured memory that is available for use. Key: managementAgent|memUsed_average
Management Agent|Memory Swap Used (KB)	Sum of the memory swapped by all powered-on virtual machines on the host. Key: managementAgent|swapUsed_average
Management Agent|Memory Swap In (KBps)	Amount of memory that is swapped in for the Service Console. Key: managementAgent|swapIn_average
Management Agent|Memory Swap Out (KBps)	Amount of memory that is swapped out for the Service Console. Key: managementAgent|swapOut_average
Management Agent|CPU Usage	CPU usage. Key: managementAgent|cpuUsage_average

Storage Path Metrics for Host Systems

Storage path metrics provide information about data storage use.

Metric	Description
StoragePath|Total Latency (ms)	Total latency in milliseconds. Key: storagePath|totalLatency
StoragePath|Total Usage (KBps)	Total latency in kilobytes per second. Key: storagePath|usage
StoragePath|Read Rate (KBps)	Rate of reading data from the virtual disk. Key: storagePath|read_average
StoragePath|Write Rate (KBps)	Rate of writing data. Key: storagePath|write_average
StoragePath|Commands per second	Average number of commands issued per second during the collection interval. Key: storagePath|commandsAveraged_average
StoragePath|Reads per second	Average number of read commands issued per second during the collection interval. Key: storagePath|numberReadAveraged_average
StoragePath|Writes per second	Average number of write commands issued per second during the collection interval. Key: storagePath|totalWriteLatency_average

Metric	Description
StoragePath|Writes per second	Average number of write commands issued per second during the collection interval. Key: storagePath|numberWriteAveraged_average
StoragePath|Read Latency (ms)	Average amount of time for a read operation by the storage adapter. Key: storagePath|totalReadLatency_average
StoragePath|Highest Latency	Highest Latency. Key: storagePath|maxTotalLatency_latest
StoragePath|Storage Path Name	Storage path name. Key: storagePath|storagePathName

Storage Adapter Metrics for Host Systems

Storage adapter metrics provide information about data storage use.

Metric	Description
Storage Adapter|Total Usage (KBps)	Total latency. Key: storageAdapter|usage
Storage Adapter|Port WWN	Port World Wide Name. Key: storageAdapter|portWWN
Storage Adapter|Commands per second	Average number of commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|commandsAveraged_average
Storage Adapter|Reads per second	Average number of read commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|numberReadAveraged_average
Storage Adapter|Writes per second	Average number of write commands issued per second by the storage adapter during the collection interval. Key: storageAdapter|numberWriteAveraged_average
Storage Adapter|Read Rate (KBps)	Rate of reading data by the storage adapter. Key: storageAdapter|read_average
* Storage Adapter|Read Latency (ms)	This metric shows the average amount of time for a read operation by the storage adapter. Use this metric to monitor the storage adapter read operation performance. A high value means that the ESXi is performing a slow storage read operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter|totalReadLatency_average
* Storage Adapter|Write Latency (ms)	This metric shows the average amount of time for a write operation by the storage adapter. Use this metric to monitor the storage adapter write performance operation. A high value means that the ESXi is performing a slow storage write operation. Total latency is the sum of kernel latency and device latency. Key: storageAdapter|totalWriteLatency_average
Storage Adapter|Write Rate (KBps)	Rate of writing data by the storage adapter. Key: storageAdapter|write_average
Storage Adapter|Demand	Demand. Key: storageAdapter|demand

Metric	Description
Storage Adapter|Highest Latency	Highest Latency. Key: torageAdapter|maxTotalLatency_latest
Storage Adapter|Outstanding Requests	Outstanding Requests. Key: storageAdapter|outstandingIOs_average
Storage Adapter|Queue Depth	Queue Depth. Key: storageAdapter|queueDepth_average
Storage Adapter|Queue Command Latency (ms)	The average time spent in the ESX Server VM Kernel queue per command. Key: storageAdapter|queueLatency_average
Storage Adapter|Queued	Queued. Key: storageAdapter|queued_average

Storage Metrics for Host Systems

Storage metrics provide information about storage use.

Metric	Description
Storage|Commands per second	Average number of commands issued per second during the collection interval. Key: storage|commandsAveraged_average
Storage|Read Latency (ms)	Average amount of time for a read operation in milliseconds. Key: storage|totalReadLatency_average
Storage|Read Rate (KBps)	Read throughput rate in kilobytes. Key: storage|read_average
Storage|Reads per second	Average number of read commands issued per second during the collection interval. Key: storage|numberReadAveraged_average
Storage|Total Latency (ms)	Total latency in milliseconds. Key: storage|totalLatency_average
Storage|Total Usage (KBps)	Total throughput rate in kilobytes per second. Key: storage|usage_average
Storage|Write Latency (ms)	Average amount of time for a write operation in milliseconds. Key: storage|totalWriteLatency_average
Storage|Write Rate (KBps)	Write throughput rate in kilobytes per second. Key: storage|write_average
Storage|Writes per second	Average number of write commands issued per second during the collection interval. Key: storage|numberWriteAveraged_average

Sensor Metrics for Host Systems

Sensor metrics provide information about host system cooling.

Metric	Description
Sensor|Fan|Speed (%)	Percent fan speed. Key: Sensor|fan|currentValue
Sensor|Fan|Health State	Fan health state. Key: Sensor|fan|healthState
Sensor|Temperature|Temp C	Fan temperature in centigrade. Key: Sensor|temperature|currentValue
Sensor|Temperature|Health State	Fan health state. Key: Sensor|temperature|healthState

Power Metrics for Host Systems

Power metrics provide information about host system power use.

Metric	Description
Power|Energy (Joule)	Host power use in joules. Key: power|energy_summation
Power|Power (Watt)	Host power use in watts. Key: power|power_average
Power|Power Cap (Watt)	Host power capacity in watts. Key: power|powerCap_average

Disk Space Metrics for Host Systems

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared disk space in gigabytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disks	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Shared Used (GB)	Used shared disk space in gigabytes. Key: diskspace|shared
Diskspace|Snapshot	Disk space used by snapshots in gigabytes. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Disk space used by virtual disks in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Disk space used by virtual machines in gigabytes. Key: diskspace|used
Diskspace|tTotal disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk spacey	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned .

Summary Metrics for Host Systems

Summary metrics provide information about overall host system performance.

Metric	Description
* Summary|Number of Running VMs	This metric shows the number of VMs running on the host during the last metric collection time. Large spikes of running VMs might be a reason for CPU or memory spikes as more resources will be used in the host. Number of Running VMs gives you a good indicator of how many requests the ESXi host must juggle. This excludes powered off VMs as they do not impact ESXi performance. A change in this number in your environment can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, as all the VMs will go down (or be relocated by HA) if the ESXi crashes. Look for any correlation between spikes in the number of running VMs and spikes in other metrics such as CPU Contention/Memory Contention. Key: summary|number_running_vms
Summary|Maximum Number of VMs	Maximum number of virtual machines Key: summary|max_number_vms
* Summary|Number of vMotions	This metric shows the number of vMotions that occurred in the host in the last X minutes. The number of vMotions is a good indicator of stability. In a healthy environment, this number should be stable and relatively low. Look for correlation between vMotions and spikes in other metrics such as CPU/Memory contention. The vMotion should not create any spikes, however, the VMs moved into the host might create spikes in memory usage, contention and CPU demand and contention. Key: summary|number_vmotion
Summary|Total Number of Datastores	Total Number of Datastores. Key: summary|total_number_datastores
Summary|Number of VCPUs on Powered On VMs	Total number of VCPUs of Virtual Machines that are powered on. Key: summary|number_running_vcpus
Summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
Summary|Workload Indicator (%)	Percent workload indicator. Key: summary|workload_indicator

HBR Metrics for Host Systems

Host-based replication (HBR) metrics provide information about vSphere replication.

Metric	Description
HBR|Replication Data Received Rate	Replication Data Received Rate. Key: hbr|hbrNetRx_average
HBR|Replication Data Transmitted Rate	Replication Data Transmitted Rate. Key: hbr|hbrNetTx_average
HBR|Replicated VM Count	Number of replicated virtual machines. Key: hbr|hbrNumVms_average

Cluster Compute Resource Metrics

vRealize Operations Manager collects configuration, storage, disk space, CPU use, disk, memory, network, power, and summary metrics for cluster compute resources.

Cluster Compute Resource metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the clusters in your environment.

Configuration Metrics for Cluster Compute Resources

Configuration metrics provide information about configuration settings.

Metric	Description
Configuration|Failover Level	DAS configuration failover level. Key: configuration|dasconfig|failoverLevel
Configuration|Active Admission Control Policy	DAS configuration active admission control policy. Key: configuration|dasconfig|activeAdministrationControlPolicy
Configuration|CPU Failover Resources Percent	Percent CPU failover resources for DAS configuration admission control policy. Key: configuration|dasconfig|adminissionControlPolicy|cpuFailoverResourcesPercent
Configuration|Memory Failover Resources Percent	Percent memory failover resources for DAS configuration admission control policy. Key: configuration|dasconfig|adminissionControlPolicy|memoryFailoverResourcesPercent

Storage Metrics for Cluster Compute Resources

Storage metrics provide information about storage use.

Metric	Description
Storage|Total Usage	Total throughput rate in kilobytes per second. Key: storage|usage_average

Disk Space Metrics for Cluster Compute Resources

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Virtual machine used (GB)	Space used by virtual machine files in gigabytes. Key: diskspace|used
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity

Metric	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|Not Shared (GB)	Space used by VMs that is not shared. Key: diskspace|notshared

CPU Usage Metrics for Cluster Compute Resources

CPU usage metrics provide information about CPU use.

Metric	Description
CPU|Capacity Usage	This metric shows the percentage of the capacity used. Key: cpu|capacity_usagepct_average
* CPU|CPU Contention (%)	This metric is an indicator of the overall contention for CPU resources that occurs across the workloads in the cluster. When contention occurs, it means that some of the virtual machines are not immediately getting the CPU resources they are requesting. Use this metric to identify when a lack of CPU resources might be causing performance issues in the cluster. This metric is the sum of the CPU contention across all hosts in the cluster averaged over two times the number of physical CPUs in the cluster to account for hyper-threading. CPU contention takes into account: ■ CPU Ready ■ CPU Co-stop ■ Power management ■ Hyper threading This metric is more accurate than CPU Ready since it takes into account CPU Co-stop and Hyper threading. When using this metric, the number should be lower than the performance you expect. If you expect performance at 10%, then the number should be lower than 10%. Since this value is averaged across all hosts in the cluster, you may find that some hosts have a higher CPU contention while others are lower. In order to ensure that vSphere spreads out the running workloads across hosts, consider enabling a fully-automated DRS in the cluster. Key: cpu|capacity_contentionPct
* CPU|Demand (%)	This metric is an indicator of the overall demand for CPU resources by the workloads in the cluster. It shows the percentage of CPU resources that all the virtual machines would use if there were no CPU contention or CPU limits set. It represents the average active CPU load in the past five minutes. Key: cpu|demandPct
CPU|Demand (MHz)	Demand in megahertz. Key: cpu|demandmhz

Metric	Description
CPU|IO Wait	IO wait time in milliseconds. Key: cpu|iowait
CPU|Number of CPU Sockets	Number of CPU sockets. Key: cpu|numpackages
CPU|Overall CPU Contention	Overall CPU contention in milliseconds. Key: cpu|capacity_contention
CPU|Host Provisioned Capacity	Provisioned CPU capacity in megahertz. Key: cpu|capacity_provisioned
CPU|Provisioned vCPUs	Number of provisioned CPU cores. Key: cpu|corecount_provisioned
CPU|Reserved Capacity	The sum of the reservation properties of the (immediate) children of the host's root resource pool in megahertz. Key: cpu|reservedCapacity_average
CPU|Wait	CPU time spent on idle state in milliseconds. Key: cpu|wait
CPU|Usage (MHz)	Average CPU use in megahertz. Key: cpu|usagemhz_average
CPU|Total Capacity	Total CPU capacity in megahertz. Key: cpu|totalCapacity _average
CPU|Demand	CPU Demand. Key: cpu|demand_average
CPU|Overhead	Amount of CPU overhead. Key: cpu|overhead_average
CPU|Demand without overhead	Value of demand excluding any overhead. Key: cpu|demand_without_overhead
CPU|Provisioned Capacity	Provisioned Capacity (MHz). Key: cpu|vm_capacity_provisioned
CPU|Number of hosts stressed	Number of hosts stressed. Key: cpu|num_hosts_stressed
CPU|Stress Balance Factor	Stress Balance Factor. Key: cpu|stress_balance_factor
CPU|Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining. Key: cpu|min_host_capacity_remaining
CPU|Workload Balance Factor	Workload Balance Factor. Key: cpu|workload_balance_factor
CPU|Highest Provider Workload	Highest Provider Workload. Key: cpu|max_host_workload
CPU|Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: cpu|host_workload_disparity
CPU|Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: cpu|host_stress_disparity

Disk Metrics for Cluster Compute Resources

Disk metrics provide information about disk use.

Metric	Description
Disk|Commands per second	Average number of commands issued per second during the collection interval. Key: disk|commandsAveraged_average
Disk|Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Command Latency and Physical Device Command Latency metrics. Key: disk|totalLatency_average
Disk|Disk Read Latency	Average amount of time for a read operation from the virtual disk. The total latency is the sum of Kernel latency and device latency. Key: disk|totalReadLatency_average
Disk|Disk Write Latency	The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency. Key: disk|totalWriteLatency_averag
Disk|Read Rate (KBps)	Number of times data was read from the disk in the defined interval. Key: disk|numberRead_summation
Disk|Reads per second	Average number of read commands issued per second during the collection interval. Key: disk|numberReadAveraged_averag
Disk|Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine. Key: disk|usage_average
Disk|Write Rate (KBps)	Number of times data was written to disk during the collection interval. Key: disk|numberWrite_summation
Disk|Writes per second	Average number of write commands issued per second during the collection interval. Key: disk|numberWriteAveraged_average
Disk|Read Requests	Amount of data read from the disk during the collection interval. Key: disk|read_average
Disk|Write Requests	Amount of data written to the disk during the collection interval. Key: disk|write_average
Disk|Commands Issued	Number of disk commands issued during the collection interval. Key: disk|commands_summation
Disk|Total Queued Outstanding operations	Sum of queued operation and outstanding operations. Key: disk|sum_queued_oio
Disk|Max Observed OIO	Max observed outstanding IO for a disk. Key: disk|max_observed

Memory Metrics for Cluster Compute Resources

Memory metrics provide information about memory use and allocation.

Metric	Description
Mem|Active Write (KB)	Active writes in kilobytes. Key: mem|activewrite_average
Mem|Compressed (KB)	Average compression in kilobytes. Key: mem|compressed_average
Mem|Compression Rate (KBps)	Average compression rate in kilobytes. Key: mem|compressionRate_average
Mem|Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory. Key: mem|consumed_average
* Mem|Contention (%)	This metric is an indicator of the overall contention for memory resources that occurs across the workloads in the cluster. When contention occurs, it means that some portion of the VMs are not immediately getting the memory resources that they are requesting. Use this metric to identify when lack of memory resources might be causing performance issues in the cluster. Key: mem|host_contentionPct
Mem|Contention (KB)	Contention in kilobytes. Key: mem|host_contention
Mem|Decompression Rate (KBps	Decompression rate in kilobytes. Key: mem|decompressionRate_average
Mem|Granted (KB)	Amount of memory available for use. Key: mem|granted_average
Mem|Guest Active (KB)	Amount of memory that is actively used. Key: mem|active_average
Mem|Heap (KB)	Amount of memory allocated for heap. Key: mem|heap_average
Mem|Heap Free (KB)	Free space in the heap. Key: mem|heapfree_average
* Mem|Balloon	This metric shows the amount of memory currently used by the virtual machine memory control. It is only defined at the VM level. Key: mem|vmmemctl_average
Mem|VM Overhead (KB)	Memory overhead reported by host. Key: mem|overhead_average
Mem|Provisioned Memory (KB)	Provisioned memory in kilobytes. Key: mem|host_provisioned
Mem|Reserved Capacity (KB)	Reserved capacity in kilobytes. Key: mem|reservedCapacity_average
Mem|Shared (KB)	Amount of shared memory. Key: mem|shared_average
Mem|Shared Common (KB)	Amount of shared common memory. Key: mem|sharedcommon_average
Mem|Swap In (KB)	Amount of memory that is swapped in for the service console. Key: mem|swapin_average
Mem|Swap In Rate (KBps)	Rate at which memory is swapped from disk into active memory during the interval. Key: mem|swapinRate_average

Metric	Description
Mem|Swap Out (KB)	Amount of memory that is swapped out for the service console. Key: mem|swapout_average
Mem|Swap Out Rate (KBps)	Rate at which memory is being swapped from active memory into disk during the current interval. Key: mem|swapoutRate_average
Mem|Swap Used (KB)	Amount of memory used for swap space. Key: mem|swapused_average
Mem|Total Capacity (KB)	Total capacity in kilobytes. Key: mem|totalCapacity_average
Mem|Reserved (KB)	Amount of unreserved memory. Key: mem|unreserved_average
Mem|Usable Memory (KB)	Usable memory in kilobytes. Key: mem|host_usable
Mem|Usage/Usable	Percent memory used. Key: mem|host_usagePct
Mem|Host Usage (KB)	Memory use in kilobytes. Key: mem|host_usage
Mem|Machine Demand	Memory Machine Demand in KB. Key: mem|host_demand
Mem|ESX System Usage	Memory usage by the VMkernel and ESX user-level services. Key: mem|host_systemUsage
* Mem|Usage (%)	This metric shows the portion of the total memory in all hosts in the cluster that is being used. This metric is the sum of memory consumed across all hosts in the cluster divided by the sum of physical memory across all hosts in the cluster. ∑ memory consumed on all hosts ---------- X 100% ∑ physical memory on all hosts
Mem|Usage (KB)	Memory currently in use as a percentage of total available memory. Key: mem|usage_average
Mem|VM kernel Usage (KB)	Amount of memory that the VM kernel uses. Key: mem|sysUsage_average
Mem|Zero (KB)	Amount of memory that is all 0. Key: mem|zero_average
Mem|Number of Hosts Stressed	Number of hosts stressed. Key: mem|num_hosts_stressed
Mem|Stress Balance Factor	Stress balance factor. Key: mem|stress_balance_factor
Mem|Lowest Provider Capacity Remaining	Lowest provider capacity remaining. Key: mem|min_host_capacity_remaining
Mem|Workload Balance Factor	Workload balance factor. Key: mem|workload_balance_factor
Mem|Highest Provider Workload	Highest provider workload. Key: mem|max_host_workload

Metric	Description
Mem|Host workload Max-Min Disparity	Difference of Max and Min host workload in the container. Key: mem|host_workload_disparity
Mem|Host stress Max-Min Disparity	Difference of Max and Min host stress in the container. Key: mem|host_stress_disparity

Network Metrics for Cluster Compute Resources

Network metrics provide information about network performance.

Metric	Description
Net|Data Receive Rate (KBps)	Average amount of data received per second. Key: net|received_average
Net|Data Transmit Rate (KBps)	Average amount of data transmitted per second. Key: net|transmitted_average
Net|Packets Dropped	Number of packets dropped in the performance interval. Key: net|dropped
Net|Packets Dropped (%)	Percentage of packets dropped. Key: net|droppedPct
Net|Packets Received	Number of packets received in the performance interval. Key: net|packetsRx_summation
Net|Packets Transmitted	Number of packets transmitted in the performance interval. Key: net|packetsTx_summation
Net|Received Packets Dropped	Number of received packets dropped in the performance interval. Key: net|droppedRx_summation
Net|Transmitted Packets Dropped	Number of transmitted packets dropped in the performance interval. Key: net|droppedTx_summation
Net|Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. Key: net|usage_average
Net|Max Observed Throughput	Max observed rate of network throughput. Key: net|maxObservedKBps
Net|Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput. Key: net|maxObserved_Tx_KBps
Net|Max Observed Received Throughput	Max observed received rate of network throughput. Key: net|maxObserved_Rx_KBps

Datastore Metrics for Cluster Compute Resources

Datastore metrics provide information about Datastore use.

Metric	Description
Datastore|Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read

Metric	Description
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write
Datastore|Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations. Key: datastore|maxObserved_OIO
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
Datastore|Reads per second	Average number of read commands issued per second during the collection interval. Key: datastore|numberReadAveraged_average
Datastore|Writes per second	Average number of write commands issued per second during the collection interval. Key: datastore|numberWriteAveraged_average
Datastore|Read Rate	Amount of data read in the performance interval. Key: datastore|read_average
Datastore|Write Rate	Amount of data written to disk in the performance interval. Key: datastore|write_average

Cluster Services Metrics for Cluster Compute Resources

Cluster Services metrics provide information about cluster services.

Metric	Description
ClusterServices|Effective CPU Resources (MHz)	VMware DRS effective CPU resources available. Key: clusterServices|effectivecpu_average
ClusterServices|Effective Memory Resources (KB)	VMware DRS effective memory resources available. Key: clusterServices|effectivemem_average

Power Metrics for Cluster Compute Resources

Power metrics provide information about power use.

Metric	Description
Power|Energy (Joule)	Energy use in joules. Key: power|energy_summation
Power|Power (Watt)	Average power use in watts. Key: power|power_average
Power|Power Cap (Watt)	Average power capacity in watts. Key: power|powerCap_average

Summary Metrics for Cluster Compute Resources

Summary metrics provide information about overall performance.

Metric	Description
Summary|Number of Running Hosts	Number of running hosts. Key: summary|number_running_hosts
* Summary|Number of Running VMs	This metric shows the total number of VMs running on all hosts in the cluster. Key: summary|number_running_vms
* Summary|Number of vMotions	This metric shows the number of vMotions that occurred during the last collection cycle. When using this metric, look for a low number which indicates that the cluster is able to serve its VMs. A vMotion can impact VM performance during the stun time. Key: summary|number_vmotion
Summary|Total Number of Hosts	Total number of hosts. Key: summary|total_number_hosts
Summary|Total Number of VMs	Total number of virtual machines. Key: summary|total_number_vms
Summary|Maximum Number of VMs	Maximum Number of virtual machines. Key: summary|max_number_vms
Summary|Workload Indicator	Percent workload indicator. Key: summary|workload_indicator
Summary|Total Number of Datastores	Total number of datastores. Key: summary|total_number_datastores
Summary|Number of VCPUs on Powered On VMs	Number of virtual CPUs on powered-on virtual machines. Key: summary|number_running_vcpus
Summary|Average Running VM Count per Running Host	Average number of running virtual machines per running host. Key: summary|avg_vm_density
Summary|Average Provisioned Capacity (MHz) per Running VM	Average provisioned capacity, in megahertz, per running virtual machine. Key: summary|avg_vm_cpu
Summary|Average Provisioned Memory (KB) per Running VM	Average provisioned memory, in kilobytes. per running virtual machine. Key: summary|avg_vm_mem

Resource Pool Metrics

vRealize Operations Manager collects configuration, CPU usage, memory, and summary metrics for resource pool objects.

Resource Pool metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

Configuration Metrics for Resource Pools

Configuration metrics provide information about memory and CPU allocation configuration.

Table 1-14. Configuration Metrics for Resource Pools

Metric Key	Metric Name	Description
config|mem_alloc_reservation	Memory Allocation Reservation	Memory Allocation Reservation.

CPU Usage Metrics for Resource Pools

CPU usage metrics provide information about CPU use.

Table 1-15. CPU Usage Metrics for Resource Pools

Metric Key	Metric Name	Description
cpu|capacity_demandEntitlementPct	Capacity Demand Entitlement (%)	CPU Capacity Demand Entitlement Percentage.
cpu|capacity_entitlement	Capacity entitlement (MHz)	CPU Capacity Entitlement.
cpu|capacity_contentionPct	CPU Contention (%)	CPU capacity contention.
cpu|demandmhz	Demand (MHz)	CPU demand in megahertz.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|usagemhz_average	Usage	Average CPU use in megahertz.
cpu|effective_limit	Effective limit	CPU effective limit.
cpu|reservation_used	Reservation Used	CPU reservation used.
cpu|estimated_entitlement	Estimated entitlement	CPU estimated entitlement.
cpu|dynamic_entitlement	Dynamic entitlement	CPU dynamic entitlement.
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead

Memory Metrics for Resource Pools

Memory metrics provide information about memory use and allocation.

Table 1-16. Memory Metrics for Resource Pools

Metric Key	Metric Name	Description
mem|vmmemctl_average	Balloon (KB)	Amount of memory currently used by the virtual machine memory control.
mem|compressionRate_average	Compression Rate (KBps)	Compression rate in kilobytes per second.
mem|consumed_average	Consumed (KB)	Amount of host memory consumed by the virtual machine for guest memory.
mem|host_contentionPct	Contention (%)	Macine contention percentage.
mem|guest_usage	Guest usage	Guest memory entitlement.
mem|guest_demand	Guest demand	Guest memory entitlement.
mem|host_contention	Contention (KB)	Machine contention in kilobytes.
mem|decompressionRate_average	Decompression Rate (KBps)	Decompression rate in kilobytes per second.
mem|granted_average	Granted (KB)	Average of memory available for use.
mem|active_average	Guest Active (KB)	Amount of memory that is actively used.
mem|overhead_average	VM Overhead (KB)	Memory overhead reported by host.

Table 1-16. Memory Metrics for Resource Pools (Continued)

Metric Key	Metric Name	Description
mem|shared_average	Shared (KB)	Amount of shared memory.
mem|reservation_used	Reservation Used	Memory Reservation Used.
mem|dynamic_entitlement	Dynamic Entitlement	Memory Dynamic Entitlement.
mem|effective_limit	Effective Limit	Memory Effective Limit.
mem|swapinRate_average	swapinRate_average	Rate at which memory is swapped from disk into active memory during the interval.
mem|swapoutRate_average	swapoutRate_average	Rate at which memory is being swapped from active memory to disk during the current interval.
mem|swapped_average	Swapped (KB)	Amount of unreserved memory.
mem|usage_average	Usage (%)	Memory currently in use as a percentage of total available memory.
mem|zero_average	Zero (KB)	Amount of memory that is all zero.
mem|zipped_latest	Zipped (KB)	Latest zipped memory in kilobytes.
mem|swapin_average	Swap In (KB)	Amount of memory swapped in kilobytes.
mem|swapout_average	Swap Out (KB)	Amount of memory swapped out in kilobytes.
mem|swapused_average	Swap Used (KB)	Amount of memory used for swap space in kilobytes.
mem|guest_provisioned	Guest Configured Memory (KB)	Guest configured memory in kilobytes.

Summary Metrics for Resource Pools

Summary metrics provide information about overall performance.

Table 1-17. Summary Metrics for Resource Pools

Metric Key	Metric Name	Description
summary|number_running_vms	Number of Running VMs	Number of running virtual machines.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|iowait	IO Wait (ms)	IO wait time in milliseconds.

Datacenter Metrics

vRealize Operations Manager collects CPU usage, disk, memory, network, storage, disk space, and summary metrics for datacenter objects.

Datacenter metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

CPU Usage Metrics for Datacenters

CPU usage metrics provide information about CPU use.

Table 1-18. CPU Usage Metrics for Datacenters

Metric Key	Metric Name	Description
cpu|capacity_usagepct_average	Capacity Usage (%)	Percent capacity used.
cpu|capacity_contentionPct	CPU Contention (%)	CPU capacity contention.
cpu|demandPct	Demand (%)	CPU demand percentage.
cpu|demandmhz	Demand	Demand in megahertz.
cpu|demand_average	Demand (MHz)	CPU Demand.
cpu|overhead_average	Overhead (KB)	Amount of CPU overhead.
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|wait	Total Wait	CPU time spent on idle state.
cpu|numpackages	Number of CPU Sockets	Number of CPU sockets.
cpu|capacity_contention	Overall CPU Contention (ms)	Overall CPU contention in milliseconds.
cpu|capacity_provisioned	Host Provisioned Capacity (MHz)	Host provisioned capacity in megahertz.
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s).
cpu|reservedCapacity_average	Reserved Capacity (MHz)	The sum of the reservation properties of the (immediate) children of the host's root resource pool.
cpu|usagemhz_average	Usage	Average CPU usage in megahertz.
cpu|iowait	IO Wait	IO wait time in milliseconds.
cpu|vm_capacity_provisioned	Provisioned Capacity	Provisioned Capacity.
cpu|stress_balance_factor	Stress Balance Factor	Stress Balance Factor.
cpu|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining.
cpu|workload_balance_factor	Workload Balance Factor	Workload Balance Factor.
cpu|max_host_workload	Highest Provider Workload	Highest Provider Workload.
cpu|host_workload_disparity	Host workload Max-Min Disparity	Difference of Max and Min host workload in the container.
cpu|host_stress_disparity	Host stress Max-Min Disparity	Difference of Max and Min host stress in the container.

Disk Metrics for Datacenters

Disk metrics provide information about disk use.

Table 1-19. Disk Metrics for Datacenters

Metric Key	Metric Name	Description
disk|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection interval.
disk|totalLatency_average	Disk Command Latency (ms)	Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Disk Command Latency and Physical Device Command Latency metrics.

Table 1-19. Disk Metrics for Datacenters (Continued)

Metric Key	Metric Name	Description
disk|usage_average	Usage Rate (KBps)	Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.
disk|sum_queued_oio	Total queued outstanding operations	Sum of queued operations and outstanding operations.
disk|max_observed	Max observed OIO	Max observed IO for a disk.

Memory Metrics for Datacenters

Memory metrics provide information about memory use and allocation.

Table 1-20. Memory Metrics for Datacenters

Metric Key	Metric Name	Description
mem|host_contentionPct	Contention (%)	Machine Contention Percentage.
mem|host_demand	Machine Demand (KB)	Memory machine demand in kilobytes.
mem|host_systemUsage	ESX System Usage	Memory usage by the VM kernel and ESX user-level services.
mem|host_provisioned	Provisioned Memory (KB)	Provisioned host memory in kilobytes.
mem|reservedCapacity_average	Reserved Capacity (KB)	Reserved memory capacity in kilobytes.
mem|host_usable	Usable Memory (KB)	Usable host memory in kilobytes.
mem|host_usage	Host Usage	Host memory use in kilobytes.
mem|host_usagePct	Usage/Usable (%)	Percent host memory used.
mem|overhead_average	VM Overhead	Memory overhead reported by host.
mem|stress_balance_factor	Stress Balance Factor	Stress Balance Factor.
mem|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest Provider Capacity Remaining.
mem|workload_balance_factor	Workload Balance Factor	Workload Balance Factor.
mem|max_host_workload	Highest Provider Workload	Highest Provider Workload.
mem|host_workload_disparity	Host workload Max-Min Disparity	Difference of Max and Min host workload in the container.
mem|host_stress_disparity	Host stress Max-Min Disparity	Difference of Max and Min host stress in the container.

Network Metrics for Datacenters

Network metrics provide information about network performance.

Table 1-21. Network Metrics for Datacenters

Metric Key	Metric Name	Description
net|droppedPct	Packets Dropped	Percentage of packets dropped.
net|maxObservedKBps	Max Observed Throughput	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput	Max observed received rate of network throughput.

Table 1-21. Network Metrics for Datacenters (Continued)

Metric Key	Metric Name	Description
net|transmitted_average	Data Transmit Rate	Average amount of data transmitted per second.
net|received_average	Data Receive Rate	Average amount of data received per second.
net|usage_average	Usage Rate (KBps)	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.

Storage Metrics for Datacenters

Storage metrics provide information about storage use.

Table 1-22. Storage Metrics for Datacenters

Metric Key	Metric Name	Description
storage|usage_average	Total Usage	Total throughput rate.

Datastore Metrics for Datacenters

Datastore metrics provide information about Datastore use.

Table 1-23. Datastore Metrics for Datacenters

Metric Key	Metric Name	Description
datastore|maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore|maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Max Observed Number of Outstanding IO Operations.
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore|numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore|numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.

Disk Space Metrics for Datacenters

Disk space metrics provide information about disk use.

Table 1-24. Disk Space Metrics for Datacenters

Metric Key	Metric Name	Description
diskspace|used	Virtual machine used	Used virtual machine disk space in gigabytes.
diskspace|total_usage	Total disk space used	Total disk space used on all datastores visible to this object.
diskspace|total_capacity	Total disk space	Total disk space on all datastores visible to this object.
diskspace|total_provisioned	Total provisioned disk space	Total provisioned disk space on all datastores visible to this object.
diskspace|notshared	Not Shared (GB)	Unshared disk space in gigabytes.
diskspace|shared	Shared Used (GB)	Shared disk space in gigabytes.
diskspace|snapshot	Snapshot Space (GB)	Snapshot disk space in gigabytes.
diskspace|diskused	Virtual Disk Used (GB)	Used virtual disk space in gigabytes.
diskspace|numvmdisk	Number of Virtual Disks	Number of Virtual Disks.

Summary Metrics for Datacenters

Summary metrics provide information about overall performance.

Table 1-25. Summary Metrics for Datacenters

Metric Key	Metric Name	Description
summary|number_running_hosts	Number of Running Hosts	Number of hosts that are ON.
summary|number_running_vms	Number of Running VMs	Number of running virtual machines.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.
summary|total_number_clusters	Total Number of Clusters	Total number of clusters.
summary|total_number_hosts	Total Number of Hosts	Total number of hosts.
summary|total_number_vms	Total Number of VMs	Total number of virtual machines.
summary|total_number_datastores	Total Number of Datastores	Total number of datastores.
summary|number_running_vcpus	Number of VCPUs on Powered On VMs	Total number of VCPUs of virtual machines that are powered on.
summary|workload_indicator	Workload Indicator	Workload indicator.
summary|avg_vm_density	Average Running VM Count per Running Host	Average number of running virtual machines per running host.

Custom Datacenter Metrics

vRealize Operations Manager collects CPU usage, memory, summary, network, and datastore metrics for custom datacenter objects.

Custom datacenter metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

CPU Usage Metrics for Custom Datacenters

CPU usage metrics provide information about CPU use.

Table 1-26. CPU Usage Metrics for Custom Datacenters

Metric Key	Metric Name	Description
cpu|capacity_provisioned	Host Provisioned Capacity	Host provisioned capacity (MHz).
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s).
cpu|demand_without_overhead	Demand without overhead	Value of demand excluding any overhead.
cpu|num_hosts_stressed	Number of hosts stressed	Number of hosts stressed.
cpu|stress_balance_factor	Stress Balance Factor	Stress balance factor.
cpu|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest provider capacity remaining.
cpu|workload_balance_factor	Workload Balance Factor	Workload balance factor.
cpu|max_host_workload	Highest Provider Workload	Highest provider workload.
cpu|host_workload_disparity	Host workload Max-Min Disparity	Host workload max-min disparity.
cpu|host_stress_disparity	Host stress Max-Min Disparity	Difference of max and min host stress in the container.

Memory Metrics for Custom Datacenters

Memory metrics provide information about memory use.

Table 1-27. Memory Metrics for Custom Datacenters

Metric Key	Metric Name	Description
mem|host_usable	Usable Memory	Usable memory.
mem|host_demand	Machine Demand	Memory machine demand in KB.
mem|num_hosts_stressed	Number of hosts stressed	Number of hosts stressed.
mem|stress_balance_factor	Stress Balance Factor	Stress balance factor.
mem|min_host_capacity_remaining	Lowest Provider Capacity Remaining	Lowest provider capacity remaining.
mem|workload_balance_factor	Workload Balance Factor	Workload balance factor.
mem|max_host_workload	Highest Provider Workload	Highest provider workload.
mem|host_workload_disparity	Host workload Max-Min Disparity	Host workload max-min disparity.
mem|host_stress_disparity		Host stress max-min disparity.

Summary Metrics for Custom Datacenters

Summary metrics provide information about overall performance.

Table 1-28. Summary Metrics for Custom Datacenters

Metric Key	Metric Name	Description
summary|number_running_vms	Number of Running VMs	Number of virtual machines that are ON.
summary|max_number_vms	Maximum Number of VMs	Maximum number of virtual machines.
summary|status	Status	Status of datacenter.

Network Metrics for Custom Datacenters

Network metrics provide information about network performance.

Table 1-29. Network Metrics for Custom Datacenters

Metric Key	Metric Name	Description
net|usage_average	Usage Rate	The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.
net|maxObserved_KBps	Max Observed Throughput	Max observed rate of network throughput.
net|maxObserved_Tx_KBps	Max Observed Transmitted Throughput	Max observed transmitted rate of network throughput.
net|maxObserved_Rx_KBps	Max Observed Received Throughput	Max observed received rate of network throughput.
net|transmitted_average	Data Transmit Rate	Average amount of data transmitted per second.
net|received_average	Data REceive Rate	Average amount of data received per second.

Datastore Metrics for Custom Datacenters

Datastore metrics provide information about datastore use.

Table 1-30. Datastore Metrics for Custom Datacenters

Metric Key	Metric Name	Description
datastore|maxObserved_NumberRead	Max Observed Reads per second	Max observed average number of read commands issued per second during the collection interval.
datastore|maxObserved_Read	Max Observed Read Rate	Max observed rate of reading data from the datastore.
datastore|maxObserved_NumberWrite	Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval.
datastore|maxObserved_Write	Max Observed Write Rate	Max observed rate of writing data from the datastore.
datastore|maxObserved_OIO	Max Observed Number of Outstanding IO Operations	Max observer number of outstanding IO operations.
datastore|demand_oio	Outstanding IO requests	OIO for datastore.
datastore|numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore|numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read rate	Amount of data read in the performance interval.
datastore|write_average	Write rate	Amount of data written to disk in the performance interval.

Storage Pod Metrics

vRealize Operations Manager collects datastore and disk space metrics for storage pod objects.

Storage Pod metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77

- "Badge Metrics," on page 80

Table 1-31. Datastore Metrics for Storage Pods

Metric Key	Metric Name	Description
datastore|numberReadAveraged_average	Reads per second	Average number of read commands issued per second during the collection interval.
datastore|numberWriteAveraged_average	Writes per second	Average number of write commands issued per second during the collection interval.
datastore|read_average	Read Rate	Amount of data read in the performance interval.
datastore|write_average	Write Rate	Amount of data written to disk in the performance interval.
datastore|usage_average	Usage Average	Usage Average.
datastore|totalReadLatency_average	Read Latency	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.
datastore|totalWriteLatency_average	Write Latency	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.
datastore|totalLatency_average	Disk Command Latency	The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.
datastore|commandsAveraged_average	Commands per second	Average number of commands issued per second during the collection interval.

Table 1-32. Diskspace Metrics for Storage Pods

Metric Key	Metric Name	Description
diskspace|disktotal	Total used	Total space used.
diskspace|freespace	Freespace	Unused space available on datastore.
diskspace|capacity	Capacity	Total capacity of datastore.
diskspace|used	Virtual Machine used	Space used by virtual machine files.
diskspace|snapshot	Snapshot Space	Space used by snapshots.

VMware Distributed Virtual Switch Metrics

vRealize Operations Manager collects network and summary metrics for VMware distributed virtual switch objects.

VMware Distributed Virtual Switch metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

Table 1-33. Network Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
network|port_statistics|rx_bytes	Total Ingress Traffic	Total ingress traffic (KBps).
network|port_statistics|tx_bytes	Total Egress Traffic	Total egress traffic (KBps).

Table 1-33. Network Metrics for VMware Distributed Virtual Switches (Continued)

Metric Key	Metric Name	Description
network|port_statistics|ucast_tx_pkts	Egress Unicast Packets per second	Egress unicast packets per second.
network|port_statistics|mcast_tx_pkts	Egress Multicast Packets per second	Egress multicast packets per second.
network|port_statistics|bcast_tx_pkts	Egress Broadcast Packets per second	Egress broadcast packets per second.
network|port_statistics|ucast_rx_pkts	Ingress Unicast Packets per second	Ingress unicast packets per second.
network|port_statistics|mcast_rx_pkts	Ingress Multicast Packets per second	Ingress multicast packets per second.
network|port_statistics|bcast_rx_pkts	Ingress Broadcast Packets per second	Ingress broadcast packets per second.
network|port_statistics|dropped_tx_pkts	Egress Dropped Packets per second	Egress dropped packets per second.
network|port_statistics|dropped_rx_pkts	Ingress Dropped Packets per second	Ingress dropped packets per second.
network|port_statistics|rx_pkts	Total Ingress Packets per second	Total ingress packets per second.
network|port_statistics|tx_pkts	Total Egress Packets per second	Total egress packets per second.
network|port_statistics|utilization	Utilization	Use (KBps).
network|port_statistics|dropped_pkts	Total Dropped Packets per second	Total dropped packets per second.
network|port_statistics|dropped_pkts_pct	Percentage of Dropped Packets	Percentage of dropped packets.
network|port_statistics|maxObserved_rx_bytes	Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps).
network|port_statistics|maxObserved_tx_bytes	Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps).
network|port_statistics|maxObserved_utilization	Max Observed Utilization (KBps)	Max observed utilization (KBps).

Table 1-34. Summary Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
summary|max_num_ports	Maximum Number of Ports	Maximum number of ports.
summary|used_num_ports	Used Number of Ports	Used number of ports.
summary|num_blocked_ports	Number of Blocked Ports	Number of blocked ports.

Table 1-35. Host Metrics for VMware Distributed Virtual Switches

Metric Key	Metric Name	Description
host|mtu_mismatch	MTU Mismatch	Maximum Transmission Unit (MTU) mismatch.
host|teaming_mismatch	Teaming Mismatch	Teaming mismatch.
host|mtu_unsupported	Unsupported MTU	Unsupported MTU.
host|vlans_unsupported	Unsupported VLANs	Unsupported VLANs.

Table 1-35. Host Metrics for VMware Distributed Virtual Switches (Continued)

Metric Key	Metric Name	Description
host|config_outofsync	Config Out Of Sync	Config Out Of Sync.
host|attached_pnics	Number of Attached pNICs	Number of attached physical NICs.

Distributed Virtual Port Group Metrics

The vCenter Adapter instance collects network and summary metrics for distributed virtual port groups.

Distributed Virtual Port Group metrics include capacity and badge metrics. See definitions in:

- "Capacity and Project-Based Metrics," on page 77
- "Badge Metrics," on page 80

Table 1-36. Network Metrics for Distributed Virtual Port Groups

Metric Key	Metric Name	Description
network|port_statistics|rx_bytes	Ingress Traffic	Ingress traffic (KBps).
network|port_statistics|tx_bytes	Egress Traffic	Egress traffic (KBps).
network|port_statistics|ucast_tx_pkts	Egress Unicast Packets per second	Egress unicast packets per second.
network|port_statistics|mcast_tx_pkts	Egress Multicast Packets per second	Egress multicast packets per second.
network|port_statistics|bcast_tx_pkts	Egress Broadcast Packets per second	Egress broadcast packets per second.
network|port_statistics|ucast_rx_pkts	Ingress Unicast Packets per second	Ingress unicast packets per second.
network|port_statistics|mcast_rx_pkts	Ingress Multicast Packets per second	Ingress multicast packets per second.
network|port_statistics|bcast_rx_pkts	Ingress Broadcast Packets per second	Ingress broadcast packets per second.
network|port_statistics|dropped_tx_pkts	Egress Dropped Packets per second	Egress dropped packets per second.
network|port_statistics|dropped_rx_pkts	Ingress Dropped Packets per second	Ingress dropped packets per second.
network|port_statistics|rx_pkts	Total Ingress Packets per second	Total Ingress packets per second.
network|port_statistics|tx_pkts	Total Egress Packets per second	Total Egress packets per second.
network|port_statistics|utilization	Utilization	Utilization (KBps).
network|port_statistics|dropped_pkts	Total Dropped Packets per second	Total dropped packets per second.
network|port_statistics|dropped_pkts_pct	Percentage of Dropped Packets	Percentage of dropped packets.
network|port_statistics|maxObserved_rx_bytes	Max Observed Ingress Traffic (KBps)	Max observed ingress traffic (KBps).
network|port_statistics|maxObserved_tx_bytes	Max Observed Egress Traffic (KBps)	Max observed egress traffic (KBps).
network|port_statistics|maxObserved_utilization	Max Observed Utilization (KBps)	Max observed utilization (KBps).

Table 1-37. Summary Metrics for Distributed Virtual Port Groups

Metric Key	Metric Name	Description
summary|max_num_ports	Maximum Number of Ports	Maximum number of ports.
summary|used_num_ports	Used Number of Ports	Used number of ports.
summary|num_blocked_ports	Number of Blocked Ports	Number of blocked ports.

Datastore Metrics

vRealize Operations Manager collects capacity, device, and summary metrics for datastore objects.

Capacity metrics can be calculated for datastore objects. See "Capacity and Project-Based Metrics," on page 77.

Metrics marked with an asterisk (*) provide the most relevant data to use when you troubleshoot the datastores in your environment.

Capacity Metrics for Datastores

Capacity metrics provide information about datastore capacity.

Metric	Description
* Capacity|Available Space (GB)	This metric shows the amount of free space that a datastore has available. Use this metric to know how much storage space is unused on the datastore. Try to avoid having too little free disk space in order to accommodate unexpected storage growth on the datastore. The exact size of the datastore is based on company policy. Key: capacity|available_space
Capacity|Data Store Capacity Contention	Datastore capacity contention. Key: capacity|contention
* Capacity|Provisioned (GB)	This metric shows the amount of storage that was allocated to the virtual machines. Use this metric to know how much storage space is currently being used on the datastore. Check the metric trend to identify spikes or abnormal growth. Key: capacity|provisioned
* Capacity|Total Capacity (GB)	This metric shows the overall size of the datastore. Use this metric to know the total capacity of the datastore. Typically the size of the datastore should not be too small. VMFS datastore size has grown over the years as virtualization matures and larger vitual machines are now onboard. Ensure the size can handle enough virtual machines to avoid datastore sprawl. A best practise is to use 5 TB for VMFS and more for vSAN. Key: capacity|total_capacity
Capacity|Used Space (GB)	This metric shows the amount of storage that is being used on the datastore. Key: capacity|used_space
Capacity|Workload (%)	Capacity workload. Key: capacity|workload
Capacity|Uncommitted Space (GB)	Uncommitted space in gigabytes. Key: capacity|uncommitted

Metric	Description
Capacity|Total Provisioned Consumer Space	Total Provisioned Consumer Space. Key: capacity|consumer_provisioned
* Capacity|Used Space (%)	This metric shows the amount of storage that is being used on the datastore. Use this metric to know the percentage of storage space being used on the datastore. When using this metric, verify that you have at least 20% of free storage. Less than this, and you may experience problems when a snapshot is not deleted. If you have more than 50% free storage space, you are not utilizing your storage in the best possible way. Key: capacity|usedSpacePct

Device Metrics for Datastores

Device metrics provide information about device performance.

Metric	Description
Devices|Bus Resets	This metric shows the number of bus resets in the performance interval. Key: devices|busResets_summation
Devices|Commands Aborted	This metric shows the number of disk commands aborted in the performance interval. Key: devices|commandsAborted_summation
Devices|Commands Issued	This metric shows the number of disk commands issued in the performance interval. Key: devices|commands_summation
Devices|Disk Command Latency (ms)	This metric shows the average time taken for a command from the perspective of a guest operating system. This metric is the sum of Kernel Disk Command Latency and Physical Device Command Latency metrics. Key: devices|totalLatency_average
Devices|Disk Read Latency (ms)	This metric shows the average time taken for a read from the perspective of a guest operating system. This metric is the sum of the Kernel Disk Read Latency and Physical Device Read Latency metrics. Key: devices|totalReadLatency_averag
Devices|Disk Write Latency (ms)	This metric shows the average amount of time for a write operation to the datastore. Total latency is the sum of kernel latency and device latency. Key: devices|totalWriteLatency_average
Devices|Kernel Disk Command Latency (ms)	Average time spent in ESX Server V. Kernel per command. Key: devices|kernelLatency_average
Devices|Kernel Disk Read Latency (ms)	Average time spent in ESX host VM Kernel per read. Key: devices|kernelReadLatency_average
Devices|Kernel Disk Write Latency (ms)	Average time spent in ESX Server VM Kernel per write. Key: devices|kernelWriteLatency_average
Devices|Number of Running Hosts	Number of running hosts that are powered on. Key: devices|number_running_hosts
Devices|Number of Running VMs	Number of running virtual machines that are powered on. Key: devices|number_running_vms

Metric	Description
Devices|Physical Device Command Latency (ms)	Average time taken to complete a command from the physical device. Key: devices|deviceLatency_average
Devices|Physical Device Read Latency (ms)	Average time taken to complete a read from the physical device. Key: devices|deviceReadLatency_average
Devices|Queue Command Latency (ms)	Average time spent in the ESX Server VM Kernel queue per command. Key: devices|queueLatency_average
Devices|Queue Read Latency (ms)	Average time spent in the ESX Server VM Kernel queue per read. Key: devices|queueReadLatency_average
Devices|Queue Write Latency (ms)	Average time spent in the ESX Server VM Kernel queue per write. Key: devices|queueWriteLatency_average
Devices|Read Rate (KBps)	Amount of data read in the performance interval. Key: devices|read_average
Devices|Read Requests	Number of times data was read from the disk in the defined interval. Key: devices|numberRead_summation
Devices|Reads per second	Average number of read commands issued per second to the datastore during the collection interval. Key: devices|numberReadAveraged_average
Devices|Usage Average (KBps)	Average use in kilobytes per second. Key: devices|usage_average
Devices|Write Rate (KBps)	Amount of data written to disk in the performance interval. Key: devices|write_average
Devices|Write Requests	Number of times data was written to the disk in the defined interval. Key: devices|numberWrite_summation
Devices|Writes per second	Average number of write commands issued per second to the datastore during the collection interval. Key: devices|numberWriteAveraged_average
Devices|Commands per second	Average number of commands issued per second during the collection interval. Key: devices|commandsAveraged_average
Devices|Physical Device Write Latency (ms)	Average time taken to complete a write from the physical disk. Key: devices|deviceWriteLatency_average

Datastore Metrics for Datastores

Datastore metrics provide information about datastore use.

Metric	Description
* Datastore|Disk Command Latency (ms)	This metric shows the adjusted read and write latency at the datastore level. Adjusted means that the latency is taking into account the number of IOs. If your IO is read-dominated, the combined value is influenced by the reads. This is the average of all the VMs running in the datastore. Because it is an average, some VMs logically experience higher latency that the value shown by this metric. To see the worst latency experienced by any VM, use the Maximum VM Disk Latency metric. Use this metric to see the performance of the datastore. It is one of two key performance indicators for a datastore, the other being the Max Read Latency. The combination of Maximum and Average gives better insight into how well the datastore is coping with the demand. The number should be lower than the performance you expect. Key: datastore|totalLatency_average
Datastore|Usage Average (KBps)	Average use in kilobytes per second. Key: datastore|usage_average
Datastore|Read Latency (ms	Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency. Key: datastore|totalReadLatency_average
* Datastore|Write Latency (ms)	Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency. Key: datastore|totalWriteLatency_average
Datastore|Demand	Demand. Key: datastore|demand
Datastore|Demand Indicator	Demand Indicator. Key: datastore|demand_indicator
Datastore|Max Observed Reads per Second	Maximum observed average number of read commands issued per second during the collection interval. Key: datastore|maxObserved_NumberRead
Datastore|Max Observed Read Rate (KBps)	Max observed rate of reading data from the datastore. Key: datastore|maxObserved_Read
* Datastore|Max Observed Read Latency (ms)	This metric displays the maximum observed average amount of time for a read operation from the datastore. Use this metric to troubleshoot when the overall datastore latency is higher than expected. Look for a number that matches the overall latency. Total latency = kernel latency + device latency. Key: datastore|maxObserved_ReadLatency
Datastore|Max Observed Writes per second	Max observed average number of write commands issued per second during the collection interval. Key: datastore|maxObserved_NumberWrite
Datastore|Max Observed Write Rate (KBps)	Max observed rate of writing data from the datastore. Key: datastore|maxObserved_Write

Metric	Description
Datastore|Max Observed Write Latency (ms)	This metric displays the maximum observed average amount of time for a write operation from the datastore. Use this metric to troubleshoot when the overall datastore latency is higher than expected. Look for a number that matches the overall latency. Total latency = kernel latency + device latency. Key: datastore|maxObserved_WriteLatency
Datastore|Max Observed Number of Outstanding IO Operations	Maximum observed number of outstanding IO operations. Key: datastore|maxObserved_OIO
Datastore|Outstanding IO requests	OIO for datastore. Key: datastore|demand_oio
* Datastore|Reads per second (IOPS)	This metric displays the average number of read commands issued per second during the collection interval. Use this metric when the total IOPS is higher than expected. Drill down to see if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system. Key: datastore|numberReadAveraged_average
* Datastore|Writes per second (IOPS)	This metric displays the average number of write commands issued per second during the collection interval. Use this metric when the total IOPS is higher than expected. Drill down to see if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system. Key: datastore|numberWriteAveraged_average
Datastore|Read rate	This metric displays the amount of data read in the performance interval. Key: datastore|read_average
Datastore|Write rate	This metric displays the amount of data written to disk in the performance interval. Key: datastore|write_average

About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.
- The metric named `devices|totalLatency_average` is not supported.

As a result, vRealize Operations Manager does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

Disk Space Metrics for Datastores

Disk space metrics provide information about disk space use.

Metric	Description
Diskspace|Not Shared (GB)	Unshared space in gigabytes. Key: diskspace|notshared
Diskspace|Number of Virtual Disk	Number of virtual disks. Key: diskspace|numvmdisk
Diskspace|Provisioned Space (GB)	Provisioned space in gigabytes. Key: diskspace|provisioned
Diskspace|Shared Used (GB)	Shared used space in gigabytes. Key: diskspace|shared
* Diskspace|Snapshot Space (GB)	This metric shows the amount of space taken by snapshots on a given database. Use this metric to know how much storage space is being used by virtual machine snapshots on the datastore. Check that the snapshot is using 0 GB or minimal space. Anything over 1 GB should trigger a warning. The actual value depends on how IO intensive the virtual machines in the datastore are. Run a DT on them to detect anomaly. Clear the snapshot within 24 hours, preferably as soon as you have finished backing up, or patching. Key: diskspace|snapshot
Diskspace|Virtual Disk Used (GB)	Virtual disk used space in gigabytes. Key: diskspace|diskused
Diskspace|Virtual machine used (GB)	Virtual machine used space in gigabytes. Key: diskspace|used
Diskspace|Total disk space used	Total disk space used on all datastores visible to this object. Key: diskspace|total_usage
Diskspace|Total disk space	Total disk space on all datastores visible to this object. Key: diskspace|total_capacity
Diskspace|Total provisioned disk space	Total provisioned disk space on all datastores visible to this object. Key: diskspace|total_provisioned
Diskspace|Total used (GB)	Total used space in gigabytes. Key: diskspace|disktotal
Diskspace|Swap File Space (GB)	Swap file space in gigabytes. Key: diskspace|swap
Diskspace|Other VM Space (GB)	Other virtual machine space in gigabytes. Key: diskspace|otherused
Diskspace|Freespace (GB)	Unused space available on datastore. Key: diskspace|freespace
Diskspace|Capacity (GB)	Total capacity of datastore in gigabytes. Key: diskspace|capacity
Diskspace|Overhead	Amount of disk space that is overhead. Key: diskspace|overhead

Summary Metrics for Datastores

Summary metrics provide information about overall performance.

Metric	Description
* Summary|Total Number of Hosts	This metric shows the number of hosts that the datastore is connected to. Use this metric to know how many clusters the datastore is attached to. The number should not be too high, as a datastore should not be mounted by every host. The datastore and cluster should be paired to keep operations simple. Key: summary|total_number_hosts
* Summary|Total Number of VMs	This metric shows the number of virtual machines which save their VMDK files on the datastore. If a VM has four VMDKs stored in four datastores, the VM will be counted on each datastore. Use this metric to know how many VMs have at least one VMDK on a specific datastore. The number of VMs should be within your Concentration Risk policy. You should also expect the datastore to be well used. If only a few VMs are using the datastore, this is not considered a good use. Key: summary|total_number_vms
Summary|Maximum Number of VMs	Maximum number of virtual machines. Key: summary|max_number_vms
Summary|Workload Indicator	Workload indicator. Key: summary|workload_indicator
* Summary|Total Number of Clusters	This metric shows the number of clusters that the datastore is connected to. Key: summary|total_number_clusters

Template Metrics for Datastores

Metric	Description
Template|Virtual Machine used	Space used by virtual machine files. Key: template|used
Template|Access Time	Last access time. Key: template|accessTime

Calculated Metrics

vRealize Operations Manager calculates metrics for capacity, badges, and the health of the system. Calculated metrics apply to a subset of objects found in the `describe.xml` file that describes each adapter.

From data that the vCenter adapter collects, vRealize Operations Manager calculates metrics for objects of type:

- vSphere World
- Virtual Machine
- Host System

- Datastore

From data that the vRealize Operations Manager adapter collects, vRealize Operations Manager calculates metrics for objects of type:

- Node
- Cluster

Capacity and Project-Based Metrics

The capacity engine computes and publishes metrics that help you to plan your resource use based on consumer demand. Project-based metrics are a subset of capacity metrics that help to plan future resource use based on predicted consumer demand.

Capacity Metrics Group

For the capacity metrics group, full metric names include the name of the resource container. For example, if density metrics are computed for CPU or memory, the actual metric names appear as cpu|density or mem|density.

Only resource containers enabled for the capacity computations have relevant metrics. Not all metric types are generated for all resource containers. For example, if CPU or memory resource containers are enabled in a policy for density, but the network resource container is not, then cpu|density and mem|density metrics are calculated but network|density metrics are not.

A capacity metric definition includes resource containers that act as a consumer or a provider. For example in vSphere, the virtual machines are consumers of CPU and memory that the ESX host provides.

Table 1-38. Capacity Metrics Group

Metric Key	Metric Name	Generated for	Description
capacityRemainingUsingConsumers_average	Capacity Remaining for Average Consumer Profile	Provider	Number of average-size consumers that can fit into the capacity remaining. An average-size consumer demands 50% of total capacity.
capacityRemainingUsingConsumers_small	Capacity Remaining for Small Consumer Profile	Provider	Number of small-size consumers that can fit into the capacity remaining. A small-size consumer demands 0 - 33% of the total capacity.
capacityRemainingUsingConsumers_medium	Capacity Remaining for Medium Consumer Profile	Provider	Number of medium-size consumers that can fit into the capacity remaining. A medium-size consumer demands 33-66% of the total capacity.
capacityRemainingUsingConsumers_large	Capacity Remaining for Large Consumer Profile	Provider	Number of large-size consumers that can fit into the capacity remaining. A large-size consumer demands 66-100% of the total capacity.
capacityRemaining	Capacity Remaining (%)	Both	Percent capacity remaining in the resource container. For example, if the resource container is memory and 2 GB out of 10 GB of memory is free, the capacityRemaining = 20%.
underusedpercent	Under used (%)	Both	Percent capacity not being used.

Table 1-38. Capacity Metrics Group (Continued)

Metric Key	Metric Name	Generated for	Description
idletimepercent	Idle time (%)	Both	Percent time a resource is idle based on use over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a resource is idle for a total of 6 days out of 30 days, idletimepercent = 20%.
wasteValue	Reclaimable Capacity	Both	Amount of reclaimable capacity based on consumer demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a vSphere host is configured with 10 GB of memory but only 2 GB of memory is used on average over 30 days, then wasteValue = 8 GB.
size.recommendation	Recommended Size	Both	Capacity recommendation based on demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if consumer demand is 2 GB of memory on average over 30 days, then the capacity recommendation is 2 GB.
optimal.vConsumption.per.pConsumption	Optimal consumption ratio	Provider	Ratio of ideal resource consumption to provision based on consumer demand over time. Ideal resource consumption is when the current capacity satisfies demand. Time is a policy setting. If not set, the default period is 30 days.
vConsumption.per.pConsumption	Consumption ratio	Provider	Ratio of current resource consumption to provision based on consumer demand.
object.demand	Stress Free Demand	Both	Demand based on peak analysis of raw demand values.
object.capacity	Usable Capacity	Both	Total capacity minus buffers. Capacity buffer is a policy setting.
object.demand.percent	Effective Demand (%)	Both	Percent capacity required by effective demand.
powered.on.consumer.count	Number of powered on consumers	Both	Number of consumers that are using a resource.
base.demand	Computed Demand	Both	Demand for an object based on self or consumer demand without the peak consideration policy setting.
actual.capacity	Current size	Both	Actual capacity without buffers
wastePercent	Reclaimable Capacity (%)	Both	Percent of reclaimable capacity based on consumer demand over time. Time is a policy setting. If not set, the default period is 30 days. For example, if a vSphere host is configured with 10 GB of memory but only 2 GB of memory is used on average over 30 days, then wastePercent = 80%.

Object-level Metrics Group

Object-level metrics are calculated to track capacity use for all objects of a particular object type.

Table 1-39. Object-level Metrics Group

Metric Key	Metric Name	Description
summary|timeRemaining	Time Remaining	Time remaining before usable capacity runs out. Usable capacity excludes capacity reserved for HA and buffers.
summary|isStress	Is Stressed	Value equals 1 or a yellow badge indicates that an object is stressed. Value equals 0 or a green badge indicates that the object is not stressed. For a stress badge defined in a policy, when the stress exceeds the lowest threshold, the badge color changes from green to yellow.
summary|capacityRemainingValue	Capacity Remaining Value	Capacity remaining.
summary|oversized	Is Oversized	Indicates if an object has too much capacity configured, value of 1, or not, value of 0.
summary|idle	Is Idle	Indicates if an object is idle (value of 1) or not (value of 0).
summary|poweredOff	Powered Off	Indicates power state of an object. Value of 1 means ON and value of 0 means OFF.
summary|capacityRemainingUsingConsumers_average	Capacity Remaining (Average consumer profile)	Capacity remaining based on average consumer demand.
summary|capacityRemainingUsingConsumers_small	Capacity Remaining (Small consumer profile)	Capacity remaining based on small consumer demand.
summary|capacityRemainingUsingConsumers_medium	Capacity Remaining (Medium consumer profile)	Capacity remaining based on medium consumer demand.
summary|capacityRemainingUsingConsumers_large	Capacity Remaining (Large consumer profile)	Capacity remaining based on large consumer demand.
summary|capacityRemaining_min	Capacity Remaining (Based on instantaneous peak)	Capacity remaining based on peak demand or stress.
summary|capacity.provider.count	Number of Capacity providers	Number of capacity providers.
summary|consumer.count	Number of Capacity consumers	Number of capacity consumers.
summary|consumer.count.per.provider.count	Consumer Provider ratio	Ratio of number of consumers to number of providers.
summary|optimal.consumer.per.provider	Optimal Consumer Provider ratio	Ratio of consumer to provider that would be optimal based on consumer demand.

Project-Based Metrics

Project-based metrics are calculated for a change in resources or demand that could affect capacity at some time in the future. See *vRealize Operations Manager User Guide*. Most metrics appear with_`whatif` appended to the capacity metric name. For example, the what-if applicable metric for capacity remaining is published as `capacityRemaining_whatif`.

Badge Metrics

Badge metrics provide information for badges in the user interface. They report the health, risk, and efficiency of objects in your environment.

vRealize Operations Manager 6.x analyzes badge metric data at five-minute averages, instead of hourly. As a result, you might find that efficiency and risk badge calculations are more sensitive than in previous versions. Badge metrics continue to be published nightly.

Table 1-40. Badge Metrics

Metric Key	Metric Name	Description
badge|alert_count_critical	Alert Count Critical	Count of critical alerts on the object.
badge|alert_count_immediate	Alert Count Immediate	Count of immediate alerts on the object.
badge|alert_count_info	Alert Count Info	Count of info alerts on the object.
badge| alert_count_warning	Alert Count Warning	Count of warning alerts on the object.
badge|anomaly	Anomaly	Overall score for anomalies, on a scale of 100.
badge|capacityRemaining	Capacity Remaining	Overall score for capacity remaining, on a scale of 100.
badge|compliance	Compliance	Overall score for compliance, on a scale of 100.
badge|density	Density	Overall score for density, on a scale of 100.
badge|efficiency	Efficiency	Overall score for efficiency. The score will be one of these discrete values representing each state of the badge: Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1.
badge|efficiency_classic	Legacy Efficiency	The legacy efficiency score computed on a scale of 100 as per vCenter Operations Manager version 5.x. For backward compatibility purposes.
badge|efficiency_state	Efficiency State	Represents the state of the efficiency badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1.
badge|fault	Fault	Overall score for fault, on a scale of 100.
badge|health	Health	Overall score for health. The score will be one of these discrete values representing each state of the badge: Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1.
badge|health_classic	Legacy Health	The legacy health score computed on a scale of 100 as per vCenter Operations Manager 5.x. For backward compatibility purposes.
badge|health_state	Health State	Represents the state of health badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1

Table 1-40. Badge Metrics (Continued)

Metric Key	Metric Name	Description
badge|risk	Risk	Overall score for risk. The score will be one of these discrete values representing each state of the badge: Green - 0, Yellow - 25, Orange - 50, Red - 75, Unknown: -1.
badge|risk_classic	Legacy Risk	The legacy risk score computed on a scale of 100 as per vCenter Operations Manager 5.x. For backward compatibility purposes.
badge|risk_state	Risk State	Represents the state of risk badge with discrete values - Green: 1, Yellow: 2, Orange: 3, Red: 4, Unknown: -1.
badge|stress	Stress	Overall score of stress, on a scale of 100.
badge|timeRemaining	Time Remaining - Real Time	Overall score of real time remaining, on a scale of 100.
badge|waste	Waste	Overall score of waste, on a scale of 100.
badge|workload	Workload (%)	Overall score of workload, on a scale of 100.

System Metrics

System metrics provide information used to monitor the health of the system. They help you to identify problems in your environment.

Table 1-41. System Metrics

Metric Key	Metric Name	Description
System Attributes|health	Self - Health Score	System health score of self resource
System Attributes|all_metrics	Self - Metric Count	Number of metrics of self resource
System Attributes|ki_metrics	Self - KPI Count	Number of KPI metrics of self resource
System Attributes|active_alarms	Self - Active Anomaly Count	Number of active alarms of self resource
System Attributes|new_alarms	Self - New Anomaly Count	Number of new alarms of self resource
System Attributes|active_ki_alarms	Self - Active KPI Breach Count	Number of active KPI alarms of self resource
System Attributes|new_ki_alarms	Self - New KPI Breach Count	Number of new KPI alarms of self resource
System Attributes|total_alarms	Self - Total Anomalies	Number of total alarms of self resource
System Attributes|change_index	Self - Change Index	Change index of self resource(100 - health score)
System Attributes|child_all_metrics	Full Set - Metric Count	Number of metrics of child resources
System Attributes|child_ki_metrics	Full Set - KPI Count	Number of KPI metrics of child resources
System Attributes|child_active_alarms	Full Set - Active Anomaly Count	Number of active alarms of child resources
System Attributes|child_new_alarms	Full Set - New Anomaly Count	Number of new alarms of child resources
System Attributes|child_active_ki_alarms	Full Set - Active KPI Breach Count	Number of active KPI alarms of child resources
System Attributes|child_new_ki_alarms	Full Set - New KPI Breach Count	Number of new KPI alarms of child resources

Table 1-41. System Metrics (Continued)

Metric Key	Metric Name	Description
System Attributes|availability	Availability	Resource availability (0-down, 1-Up, -1-Unknown)
System Attributes|alert_count_critical	Alert Count Critical	Number of Critical alerts
System Attributes|alert_count_immediate	Alert Count Immediate	Number of Immediate alerts
System Attributes|alert_count_warning	Alert Count Warning	Number of Warning alerts
System Attributes|alert_count_info	Alert Count Info	Number of Info alerts

Self-Monitoring Metrics for vRealize Operations Manager

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

Analytics Metrics

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

Table 1-42. Analytics Metrics

Metric Key	Metric Name	Description
ActiveAlarms	Active DT Symptoms	Active DT Symptoms.
ActiveAlerts	Active Alerts	Active alerts.
PrimaryResourcesCount	Number of primary objects	Number of primary objects
LocalResourcesCount	Number of local objects	Number of local objects
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
LocalMetricsCount	Number of local metrics	Number of local metrics
ReceivedResourceCount	Number of received objects	Number of received objects
ReceivedMetricCount	Number of received metrics	Number of received metrics
LocalFDSize	Number of forward data entries	Number of locally stored primary and redundant entries in forward data region.
LocalPrimaryFDSize	Number of primary forward data entries	Number of locally stored primary entries in forward data region.
LocalFDAltSize	Number of alternative forward data entries	Number of locally stored primary and redundant entries in alternative forward data region.
LocalPrimaryFDAltSize	Number of alternative primary forward data entries	Number of locally stored primary entries in alternative forward data region.
CurrentHeapSize	Current heap size	Current heap size.
MaxHeapSize	Max heap size	Max heap size
CommittedMemory	Committed memory	Committed memory
CPUUsage	CPU usage	CPU usage

Table 1-42. Analytics Metrics (Continued)

Metric Key	Metric Name	Description
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 1-43. Overall Threshold Checking Metrics for the Analytics Service

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

Table 1-43. Overall Threshold Checking Metrics for the Analytics Service (Continued)

Metric Key	Metric Name	Description
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
UpdateResourceCache|ModifcationEstimateCount|TotalCount	Total	The number of estimated modifications done during each resource cache object update.
UpdateResourceCache|ModifcationEstimateCount|AvgCount	Average	Average
UpdateResourceCache|ModifcationEstimateCount|MinCount	Minimal	Minimal
UpdateResourceCache|ModifcationEstimateCount|MaxCount	Maximal	Maximal
ManageAlerts|Count	Count	The total number of times the threshold checking work items perform alert updates.

Table 1-43. Overall Threshold Checking Metrics for the Analytics Service (Continued)

Metric Key	Metric Name	Description
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 1-44. Dynamic Threshold Calculation Metrics for the Analytics Service

Metric Key	Metric Name	Description
DtDataWrite|WriteOperationCount	Write operation count	Write operation count
DtDataWrite|Duration|TotalDuration	Total	Total length of duration (ms)
DtDataWrite|Duration|AvgDuration	Average	Average duration (ms)
DtDataWrite|Duration|MinDuration	Minimum	Minimum duration (ms)
DtDataWrite|Duration|MaxDuration	Maximum	Maximum duration (ms)
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

Table 1-44. Dynamic Threshold Calculation Metrics for the Analytics Service (Continued)

Metric Key	Metric Name	Description
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 1-45. Function Call Metrics for the Analytics Service

Metric Key	Metric Name	Description
FunctionCalls|Count	Number of function calls	Number of function calls
FunctionCalls|AvgDuration	Average execution time	Average execution time
FunctionCalls|MaxDuration	Max execution time	Max execution time

Collector Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager Collector service objects.

Table 1-46. Collector Metrics

Metric Key	Metric Name	Description
ThreadpoolThreadsCount	Number of pool threads	Number of pool threads.
RejectedFDCount	Number of rejected forward data	Number of rejected forward data
RejectedFDAltCount	Number of rejected alternative forward data	Number of rejected alternative forward data
SentFDCount	Number of sent objects	Number of sent objects
SentFDAltCount	Number of alternative sent objects	Number of alternative sent objects
CurrentHeapSize	Current heap size (MB)	Current heap size.
MaxHeapsize	Max heap size (MB)	Maximum heap size.
CommittedMemory	Committed memory (MB)	Amount of committed memory.
CPUUsage	CPU usage	CPU usage.
Threads	Threads	Number of threads.
UpStatus	Up Status	Up Status

Controller Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager Controller objects.

Table 1-47. Controller Metrics

Metric Key	Metric Name	Description
RequestedMetricCount	Number of requested metrics	Number of requested metrics
ApiCallsCount	Number of API calls	Number of API calls
NewDiscoveredResourcesCount	Number of discovered objects	Number of discovered objects

FSDB Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager file system database (FSDB) objects.

Table 1-48. FSDB Metrics

Metric Key	Metric Name	Description
StoragePoolElementsCount	Number of storage work items	Number of storage work items
FsdbState	Fsdb state	Fsdb state
StoredResourcesCount	Number of stored objects	Number of stored objects
StoredMetricsCount	Number of stored metrics	Number of stored metrics

Table 1-49. Storage Thread Pool Metrics for FSDB

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

Product UI Metrics

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

Table 1-50. Product UI Metrics

Metric Key	Metric Name	Description
ActiveSessionsCount	Active sessions	Active sessions
CurrentHeapSize	Current heap size	Current heap size.
MaxHeapsize	Max heap size	Maximum heap size.
CommittedMemory	Committed memory	Amount of committed memory.
CPUUsage	CPU usage	Percent CPU use.
Threads	Threads	Number of threads.
SessionCount	Number of active sessions	Number of active sessions
SelfMonitoringQueueSize	Self Monitoring queue size	Self Monitoring queue size

Table 1-51. API Call Metrics for the Product UI

Metric Key	Metric Name	Description
APICalls|HTTPRequesterRequestCount	HTTPRequester request count	HTTPRequester request count
APICalls|AvgHTTPRequesterRequestTime	HTTPRequester average request time	HTTPRequester average request time (ms)
APICalls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
APICalls|AvgAlertRequestTime	Average alert request time	Average alert request time (ms)
APICalls|AlertRequestCount	Alert request count	Alert request count
APICalls|AvgMetricPickerRequestTime	Average metric-picker request time	Average metric-picker request time (ms)
APICalls|MetricPickerRequestCount	Metric picker request count	Metric picker request count
APICalls|HeatmapRequestCount	Heatmap request count	Heatmap request count
APICalls|AvgHeatmapRequestTime	Average HeatMap request time	Average HeatMap request time (ms)
APICalls|MashupChartRequestCount	Mashup Chart request count	Mashup Chart request count
APICalls|AvgMashupChartRequestTime	Average Mashup Chart request time	Average Mashup Chart request time (ms)
APICalls|TopNRequestCount	Top N request count	Top N request count
APICalls|AvgTopNRequestTime	Average Top N request time	Average Top N request time (ms)
APICalls|MetricChartRequestCount	Metric Chart request count	Metric Chart request count
APICalls|AvgMetricChartRequestTime	Average MetricChart request time	Average MetricChart request time (ms)

Admin UI Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager administration user interface objects.

Table 1-52. Admin UI Metrics

Metric Key	Metric Name	Description
CurrentHeapSize	Current heap size	Current heap size (MB).
MaxHeapsize	Max heap size	Maximum heap size (MB).
CommittedMemory	Committed memory	Amount of committed memory (MB) .
CPUUsage	CPU usage	CPU usage (%).
Threads	Threads	Number of threads.
SessionCount	Number of active sessions	Number of active sessions
SelfMonitoringQueueSize	Self Monitoring queue size	Self Monitoring queue size

Table 1-53. API Call Metrics for the Admin UI

Metric Key	Metric Name	Description
APICalls|HTTPRequesterRequestCount	HTTPRequester request count	HTTPRequester request count
APICalls| AvgHTTPRequesterRequestTime	HTTPRequester average request time	HTTPRequester average request time (ms)

Suite API Metrics

vRealize Operations Manager collects metrics for the VMware vRealize Operations Management Suite API objects.

Table 1-54. Suite API Metrics

Metric Key	Metric Name	Description
UsersCount	Number of users	Number of users
ActiveSessionsCount	Active sessions	Active sessions
GemfireClientReconnects	Gemfire Client Reconnects	Gemfire Client Reconnects
GemfireClientCurrentCalls	Gemfire Client Total Outstanding	Gemfire Client Total Outstanding
CurrentHeapSize	Current heap size	Current heap size (MB) .
MaxHeapsize	Max heap size	Maximum heap size (MB) .
CommittedMemory	Committed memory	Amount of committed memory (MB).
CPUUsage	CPU usage	CPU usage (%) .
CPUProcessTime	CPU process time	CPU process time (ms)
CPUProcessTimeCapacity	CPU process time capacity	CPU process time capacity (ms)
Threads	Threads	Number of threads.

Table 1-55. Gemfire Client Call Metrics for the Suite API

Metric Key	Metric Name	Description
GemfireClientCalls|TotalRequests	Total Requests	Total Requests
GemfireClientCalls|AvgResponseTime	Average Response Time	Average Response Time (ms)

Table 1-55. Gemfire Client Call Metrics for the Suite API (Continued)

Metric Key	Metric Name	Description
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 1-56. API Call Metrics for the Suite API

Metric Key	Metric Name	Description
APICalls|TotalRequests	Total Requests	Total Requests
APICalls|AvgResponseTime	Average Response Time (ms)	Average Response Time (ms)
APICalls|MinResponseTime	Minimum Response Time (ms)	Minimum Response Time (ms)
APICalls|MaxResponseTime	Maximum Response Time	Maximum Response Time
APICalls|ServerErrorResponseCount	Server Error Response Count	Server Error Response Count
APICalls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
APICalls|FailedAuthorizationCount	Failed Authorization Count	Failed Authorization Count
APICalls|RequestsPerSecond	Requests per Second	Requests per Second
APICalls|CurrentRequests	Current Requests	Current Requests
APICalls|ResponsesPerSecond	Responses per Second	Responses per Second
APICalls|RequestsCount	Requests Count	Requests Count
APICalls|ResponsesCount	Responses Count	Responses Count

Cluster and Slice Administration Metrics

vRealize Operations Manager collects metrics for vRealize Operations Manager Cluster and Slice Administration (CaSA) objects.

Table 1-57. Cluster and Slice Administration Metrics

Metric Key	Metric Name	Description
CurrentHeapSize	Current heap size	Current heap size (MB).
MaxHeapsize	Max heap size	Maximum heap size (MB).
CommittedMemory	Committed memory	Amount of committed memory (MB).
CPUUsage	CPU usage	CPU usage (%)
Threads	Threads	Number of threads.

Table 1-58. API Call Metrics for Cluster and Slice Administration

Metric Key	Metric Name	Description
API Calls|TotalRequests	Total Requests	Total Requests
API Calls|AvgResponseTime	Average Response Time	Average Response Time (ms)
API Calls|MinResponseTime	Minimum Response Time	Minimum Response Time (ms)
API Calls|MaxResponseTime	Maximum Response Time	Maximum Response Time (ms)

Table 1-58. API Call Metrics for Cluster and Slice Administration (Continued)

Metric Key	Metric Name	Description
API Calls|ServerErrorResponseCount	Server Error Response Count	Server Error Response Count
API Calls|FailedAuthenticationCount	Failed Authentication Count	Failed Authentication Count
API Calls|FailedAuthorizationCount	Minimum Response Time	Minimum Response Time (ms)

Watchdog Metrics

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

Watchdog Metrics

The watchdog metric provides the total service count.

Table 1-59. Watchdog Metrics

Metric Key	Metric Name	Description
ServiceCount	Service Count	Service Count

Service Metrics

Service metrics provide information about watchdog activity.

Table 1-60. Metrics for the vRealize Operations Manager Watchdog Service

Metric Key	Metric Name	Description
Service|Enabled	Enabled	Enabled
Service|Restarts	Restarts	Number of times the process has been unresponsive and been restarted by Watchdog.
Service|Starts	Starts	Number of times the process has been revived by Watchdog.
Service|Stops	Stops	Number of times the process has been stopped by Watchdog.

Node Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager node objects.

Metrics can be calculated for node objects. See "Calculated Metrics," on page 76.

Table 1-61. Node Metrics

Metric Key	Metric Name	Description
Component Count	Component count	The number of vRealize Operations Manager objects reporting for this node
PrimaryResourcesCount	Number of primary objects	Number of primary objects
LocalResourcesCount	Number of local objects	Number of local objects
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
LocalMetricsCount	Number of local metrics	Number of local metrics

Table 1-61. Node Metrics (Continued)

Metric Key	Metric Name	Description
PercentDBStorageAvailable	Percent disk available /storage/db	Percent disk available /storage/db
PercentLogStorageAvailable	Percent disk available /storage/log	Percent disk available /storage/log

Table 1-62. Memory Metrics for the Node

Metric Key	Metric Name	Description
mem|actualFree	Actual Free	Actual Free
mem|actualUsed	Actual Used	Actual Used
mem|free	Free	Free)
mem|used	Used	Used
mem|total	Total	Total
mem|demand_gb	Estimated memory demand	Estimated memory demand

Table 1-63. Swap Metrics for the Node

Metric Key	Metric Name	Description
swap|total	Total	Total
swap|free	Free	Free
swap|used	Used	Used
swap|pageIn	Page in	Page in
swap|pageOut	Page out	Page out

Table 1-64. Resource Limit Metrics for the Node

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

Table 1-65. Network Metrics for the Node

Metric Key	Metric Name	Description
net|allInboundTotal	All inbound connections	All inbound total
net|allOutboundTotal	All outbound connections	All outbound total
net|tcpBound	TCP bound	TCP bound
net|tcpClose	TCP state CLOSE	Number of connections in TCP CLOSE
net|tcpCloseWait	TCP state CLOSE WAIT	Number of connections in TCP state CLOSE WAIT
net|tcpClosing	TCP state CLOSING	Number of connections in TCP state CLOSING
net|tcpEstablished	TCP state ESTABLISHED	Number of connections in TCP state ESTABLISHED

Table 1-65. Network Metrics for the Node (Continued)

Metric Key	Metric Name	Description
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 1-66. Network Interface Metrics for the Node

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

Table 1-67. Disk Filesystem Metrics for the Node

Metric Key	Metric Name	Description
disk|fileSystem|total	Total	Total
disk|fileSystem|available	Available	Available
disk|fileSystem|used	Used	Used
disk|fileSystem|files	Total file nodes	Total file nodes
disk|fileSystem|filesFree	Total free file nodes	Total free file nodes
disk|fileSystem|queue	Disk queue	Disk queue
disk|fileSystem|readBytes	Read bytes	Number of bytes read

Table 1-67. Disk Filesystem Metrics for the Node (Continued)

Metric Key	Metric Name	Description
disk|fileSystem|writeBytes	Write bytes	Number of bytes written
disk|fileSystem|reads	Reads	Number of reads
disk|fileSystem|writes	Writes	Number of writes

Table 1-68. Disk Installation Metrics for the Node

Metric Key	Metric Name	Description
disk|installation|used	Used	Used
disk|installation|total	Total	Total
disk|installation|available	Available	Available

Table 1-69. Disk Database Metrics for the Node

Metric Key	Metric Name	Description
disk|db|used	Used	Used
disk|db|total	Total	Total
disk|db|available	Available	Available

Table 1-70. Disk Log Metrics for the Node

Metric Key	Metric Name	Description
disk|log|used	Used	Used
disk|log|total	Total	Total
disk|log|available	Available	Available

Table 1-71. CPU Metrics for the Node

Metric Key	Metric Name	Description
cpu|combined	Combined load	Combined load (User + Sys + Nice + Wait)
cpu|idle	Idle	Idle time fraction of total available cpu (cpu load)
cpu|irq	Irq	Interrupt time fraction of total available cpu (cpu load)
cpu|nice	Nice	Nice time fraction of total available cpu (cpu load)
cpu|softIrq	Soft Irq	Soft interrupt time fraction of total available cpu (cpu load)
cpu|stolen	Stolen	Stolen time fraction of total available cpu (cpu load)
cpu|sys	Sys	Sys time fraction of total available cpu (cpu load)
cpu|user	User (cpu load)	User time fraction of total available cpu (cpu load)
cpu|wait	Wait (cpu load)	Wait time fraction of total available cpu (cpu load)
cpu|total	Total available for a cpu	Total available for a cpu

Table 1-71. CPU Metrics for the Node (Continued)

Metric Key	Metric Name	Description
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 1-72. Device Metrics for the Node

Metric Key	Metric Name	Description
device|iops	Reads/Writes per second	Average number of read/write commands issued per second during the collection interval.
device|await	Average transaction time	Average transaction time (milliseconds).
device|iops_readMaxObserved	Maximum observed reads per second	Maximum observed reads per second.
device|iops_writeMaxObserved	Maximum observed writes per second	Maximum observed writes per second.

Table 1-73. Service Metrics for the Node

Metric Key	Metric Name	Description
service|proc|fdUsage	Total number of open file descriptors	Total number of open file descriptors.

Table 1-74. NTP Metrics for the Node

Metric Key	Metric Name	Description
ntp|serverCount	Configured server count	Configured server count
ntp|unreachableCount	Unreachable server count	Unreachable server count
ntp|unreachable	Unreachable	Is the NTP server unreachable. Value of 0 is reachable, 1 means the server was not reached or didn't respond.

Table 1-75. Heap Metrics for the Node

Metric Key	Metric Name	Description
heap|CurrentHeapSize	Current heap size	Current heap size
heap|MaxHeapSize	Max heap size	Max heap size
heap|CommittedMemory	Committed Memory	Committed Memory

Cluster Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager cluster objects including dynamic threshold calculation metrics and capacity computation metrics.

Metrics can be calculated for cluster objects. See “Calculated Metrics,” on page 76.

Cluster Metrics

Cluster metrics provide host, resource, and metric counts on the cluster.

Table 1-76. Cluster Metrics

Metric Key	Metric Name	Description
HostCount	Number of Nodes in Cluster	Number of Nodes in Cluster
PrimaryResourcesCount	Number of primary resources	Number of primary resources
LocalResourcesCount	Number of local resources	Number of local resources
PrimaryMetricsCount	Number of primary metrics	Number of primary metrics
ReceivedResourceCount	Number of received resources	Number of received resources
ReceivedMetricCount	Number of received metrics	Number of received metrics

DT Metrics

DT metrics are dynamic threshold metrics for the cluster. Non-zero values appear only if metric collection occurs while the dynamic threshold calculations are running.

Table 1-77. DT Metrics for the Cluster

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

Capacity Computation (CC) Metrics

CC metrics are capacity computation metrics for the cluster. Non-zero values appear only if metric collection occurs while the capacity computation calculations are running.

Table 1-78. CC Metrics for the Cluster

Metric Key	Metric Name	Description
cc|isRunning	Running	Running
cc|runTime	Total Run Time	Total Run Time
cc|startTime	Start time	Start time
cc|finishTime	Finish Time	Finish Time
cc|totalResourcesToProcess	Total Objects Count	Total Objects Count
cc|progress	Progress	Progress

Table 1-78. CC Metrics for the Cluster (Continued)

Metric Key	Metric Name	Description
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 1-79. Gemfire cluster Metrics for the Cluster

Metric Key	Metric Name	Description
GemfireCluster|System|AvgReads	Average reads per second	The average number of reads per second for all members
GemfireCluster|System|AvgWrites	Average writes per second	The average number of writes per second for all members
GemfireCluster|System|DiskReadsRate	Disk reads rate	The average number of disk reads per second across all distributed members
GemfireCluster|System|DiskWritesRate	Disk writes rate	The average number of disk writes per second across all distributed members
GemfireCluster|System|GarbageCollectionCount	Total garbage collection count	The total garbage collection count for all members
GemfireCluster|System|GarbageCollectionCountDelta	New garbage collection count	The new garbage collection count for all members
GemfireCluster|System|JVMPauses	JVM pause count	The number of detected JVM pauses
GemfireCluster|System|JVMPausesDelta	New JVM pause count	The number of new detected JVM pauses
GemfireCluster|System|DiskFlushAvgLatency	Disk flush average latency	Disk flush average latency (msec)
GemfireCluster|System|NumRunningFunctions	Number of running functions	The number of map-reduce jobs currently running on all members in the distributed system
GemfireCluster|System|NumClients	Number of clients	The number of connected clients
GemfireCluster|System|TotalHitCount	Total hit count	Total number of cache hits for all regions
GemfireCluster|System|TotalHitCountDelta	New hit count	Number of new cache hits for all regions
GemfireCluster|System|TotalMissCount	Total miss count	The total number of cache misses for all regions
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

Table 1-79. Gemfire cluster Metrics for the Cluster (Continued)

Metric Key	Metric Name	Description
GemfireCluster|System|Member|FreePhysicalMemory	Free physical memory	Free physical memory (MB)
GemfireCluster|System|Member|TotalPhysicalMemory	Total physical memory	Total physical memory (MB)
GemfireCluster|System|Member|CacheListenerCallsAvgLatency	Average cache listener calls latency	Average cache listener calls latency (msec)
GemfireCluster|System|Member|CacheWriterCallsAvgLatency	Average cache writer calls latency	Average cache writer calls latency (msec)
GemfireCluster|System|Member|DeserializationAvgLatency	Average deserialization latency	Average deserialization latency (msec)
GemfireCluster|System|Member|FunctionExecutionRate	Function executions per second	Function executions per second
GemfireCluster|System|Member|JVMPauses	Number of JVM pauses	Number of JVM pauses
GemfireCluster|System|Member|NumRunningFunctions	Number of running functions	Number of running functions
GemfireCluster|System|Member|PutsRate	Puts per second	Puts per second
GemfireCluster|System|Member|GetsRate	Gets per second	Gets per second
GemfireCluster|System|Member|GetsAvgLatency	Average gets latency	Average gets latency (msec)
GemfireCluster|System|Member|PutsAvgLatency	Average puts latency	Average puts latency (msec)
GemfireCluster|System|Member|SerializationAvgLatency	Average serialization latency	Average serialization latency (msec)
GemfireCluster|System|Member|Disk|DiskFlushAvgLatency	Flush average latency	Flush average latency (msec)
GemfireCluster|System|Member|Disk|DiskReadsRate	Average reads per second	Average reads per second
GemfireCluster|System|Member|Disk|DiskWritesRate	Average writes per second	Average writes per second
GemfireCluster|System|Member|Network|BytesReceivedRate	Average received bytes per second	Average received bytes per second
GemfireCluster|System|Member|Network|BytesSentRate	Average sent bytes per second	Average sent bytes per second
GemfireCluster|System|Member|JVM|GCTimeMillis	Garbage Collection time	Total amount of time spent on garbage collection
GemfireCluster|System|Member|JVM|GCTimeMillisDelta	New Garbage Collection time	New amount of time spent on garbage collection
GemfireCluster|System|Member|JVM|TotalThreads	Total threads	Total threads
GemfireCluster|System|Member|JVM|CommitedMemory	Committed Memory	Committed Memory (MB)
GemfireCluster|System|Member|JVM|MaxMemory	Max Memory	Max Memory (MB)
GemfireCluster|System|Member|JVM|UsedMemory	Used Memory	Used Memory (MB)

Table 1-79. Gemfire cluster Metrics for the Cluster (Continued)

Metric Key	Metric Name	Description
GemfireCluster|Region|SystemRegionEntryCount	Entry Count	Entry Count
GemfireCluster|Region|DestroyRate	Destroys per second	Destroys per second
GemfireCluster|Region|CreatesRate	Creates per second	Creates per second
GemfireCluster|Region|GetsRate	Gets per second	Gets per second
GemfireCluster|Region|BucketCount	Bucket count	Bucket count
GemfireCluster|Region|AvgBucketSize	Average number of entries per bucket	Average number of entries per bucket
GemfireCluster|Region|Member|ActualRedundancy	Actual redundancy	Actual redundancy
GemfireCluster|Region|Member|BucketCount	Bucket count	Bucket count
GemfireCluster|Region|Member|AvgBucketSize	Average number of entries per bucket	Average number of entries per bucket
GemfireCluster|Region|Member|CreatesRate	Creates per second	Creates per second
GemfireCluster|Region|Member|GetsRate	Gets per second	Gets per second
GemfireCluster|Region|Member|DestroyRate	Destroys per second	Destroys per second
GemfireCluster|Region|Member|MissCount	Number of misses count	Number of cache misses
GemfireCluster|Region|Member|MissCountDelta	Number of new cache misses	Number of new cache misses
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 1-80. Threshold Checking Metrics for the Cluster

Metric Key	Metric Name	Description
ThresholdChecking|ProcessedMetricCount	Number of processed metrics	Number of processed metrics
ThresholdChecking|ProcessedMetricRate	Received metric processing rate (per second)	Received metric processing rate (per second)
ThresholdChecking|ComputedMetricCount	Number of computed metrics	Number of computed metrics
ThresholdChecking|ComputedMetricRate	Computed metric processing rate (per second)	Computed metric processing rate (per second)

Memory Metrics

Memory metrics provide memory CPU use information for the cluster.

Table 1-81. Memory Metrics for the Cluster

Metric Key	Metric Name	Description
Memory|AvgFreePhysicalMemory	Average free physical memory	Average free physical memory (GB)
Memory|TotalFreePhysicalMemory	Free physical memory	Free physical memory (GB)
Memory|TotalMemory	Total Available Memory	Total Available Memory (GB)
Memory|TotalUsedMemory	Actual Used Memory	Actual Used Memory (GB)
Memory|TotalDemandMemory	Memory Demand	Memory Demand (GB)

Elastic Memory Metrics

Elastic memory metrics provide reclaimable memory CPU use information for the cluster.

Table 1-82. Memory Metrics for the Cluster

Metric Key	Metric Name	Description
ElasticMemory|TotalMemory	Total Available Memory	Total Available Memory (GB)
ElasticMemory|TotalUsedMemory	Actual Used Memory	Actual Used Memory (GB)
ElasticMemory|TotalDemandMemory	Memory Demand	Memory Demand (GB)

CPU Metrics

CPU metrics provide CPU information for the cluster.

Table 1-83. CPU Metrics for the Cluster

Metric Key	Metric Name	Description
cpu|TotalCombinedUsage	CPU Load	CPU Load
cpu|TotalAvailable	CPU Available	CPU Available
cpu|TotalAvailable_ghz	Available	Available (GHz)
cpu|TotalUsage_ghz	Used	Used (GHz)
cpu|TotalUsage	CPU usage	CPU usage (%)

Disk Metrics

Disk metrics provide available disk information for the cluster.

Table 1-84. Disk Metrics for the Cluster

Metric Key	Metric Name	Description
Disk|DatabaseStorage|AvgAvailable	Average node disk available	Average node disk available
Disk|DatabaseStorage|MinAvailable	Minimum node disk available	Minimum node disk available
Disk|DatabaseStorage|MaxAvailable	Maximum node disk available	Maximum node disk available
Disk|DatabaseStorage|TotalAvailable	Available	Available
Disk|DatabaseStorage|Total	Total	Total
Disk|DatabaseStorage|TotalUsed	Used	Used
Disk|LogStorage|AvgAvailable	Average node disk available	Average node disk available
Disk|LogStorage|MinAvailable	Minimum node disk available	Minimum node disk available
Disk|LogStorage|MaxAvailable	Maximum node disk available	Maximum node disk available

Table 1-84. Disk Metrics for the Cluster (Continued)

Metric Key	Metric Name	Description
Disk|LogStorage|TotalAvailable	Available	Available
Disk|LogStorage|Total	Total	Total
Disk|LogStorage|TotalUsed	Used	Used

Persistence Metrics

vRealize Operations Manager collects metrics for various persistence resources or service groups.

Activity Metrics

Activity metrics relate to the activity framework.

Table 1-85. 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 master database.

Table 1-86. Controller XDB Metrics for Persistence

Metric Key	Metric Name	Description
ControllerXDB|Size	Size	Size (Bytes)
ControllerXDB|TempDBSize	Temporary DB Size	Temporary DB Size (Bytes)
ControllerXDB|TotalObjectCount	Total Object Count	Total Object Count
ControllerXDB|AvgQueryDuration	Average Query Duration	Average Query Duration (ms)
ControllerXDB|MinQueryDuration	Minimum Query Duration	Minimum Query Duration (ms)
ControllerXDB|MaxQueryDuration	Maximum Query Duration	Maximum Query Duration (ms)
ControllerXDB|TotalTransactionCount	Total Transaction Count	Total Transaction Count
ControllerXDB|LockOperationErrorCount	Lock Operation Error Count	Lock Operation Error Count
ControllerXDB|DBCorruptionErrorCount	DB Corruption Error Count	DB Corruption Error Count
ControllerXDB|DBMaxSessionExceededCount	DB Maximum Sessions Exceeded Count	DB Maximum Sessions Exceeded Count
ControllerXDB|NumberWaitingForSession	Number of operations waiting for a session	Number of operations waiting for a session from the session pool
ControllerXDB|AvgWaitForSessionDuration	Average acquisition time from session pool	Average acquisition time from session pool
ControllerXDB|MinWaitForSessionDuration	Minimum acquisition time from session pool	Minimum acquisition time from session pool
ControllerXDB|MaxWaitForSessionDuration	Maximum acquisition time from session pool	Maximum acquisition time from session pool

Table 1-86. Controller XDB Metrics for Persistence (Continued)

Metric Key	Metric Name	Description
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 1-87. Alarm XDB Metrics for Persistence

Metric Key	Metric Name	Description
AlarmSQL|Size	Size (Bytes)	Size (Bytes)
AlarmSQL|AvgQueryDuration	Average Query Duration (ms)	Average Query Duration (ms)
AlarmSQL|MinQueryDuration	Minimum Query Duration (ms)	Minimum Query Duration (ms)
AlarmSQL|MaxQueryDuration	Maximum Query Duration (ms)	Maximum Query Duration (ms)
AlarmSQL|TotalTransactionCount	Total Transaction Count	Total Transaction Count
AlarmSQL|TotalAlarms	Alarm Total Object Count	Alarm Total Object Count
AlarmSQL|TotalAlerts	Alert Total Object Count	Alert Total Object Count
AlarmSQL|AlertTableSize	Alert Table Size	Alert Table Size
AlarmSQL|AlarmTableSize	Alarm Table Size	Alarm Table Size

Key Value Store Database (KVDB)

KVDB metrics relate to the persistence of storing key-value data.

Metric Key	Metric Name	Description
KVDB|AvgQueryDuration	Average Query Duration	Average Query Duration
KVDB|MinQueryDuration	Minimum Query Duration	Minimum Query Duration
KVDB|MaxQueryDuration	Maximum Query Duration	Maximum Query Duration
KVDB|TotalTransactionCount	Total Transaction Count	Total Transaction Count

Historical Inventory Service XDB Metrics

Historical inventory service metrics relate to the persistence of configuration properties and their changes.

Table 1-88. Historical XDB Metrics for Persistence

Metric Key	Metric Name	Description
HisXDB|FunctionCalls|Count HisXDB|FunctionCalls	Number of Function calls	Number of Function calls
HisXDB|FunctionCalls|AvgDuration	Average execution time	Average execution time
HisXDB|FunctionCalls|MaxDuration	Max execution time	Max execution time
HisXDB|Size	Size	Size (Bytes)
HisXDB|TempDBSize	Temporary DB Size	Temporary DB Size (Bytes)
HisXDB|TotalObjectCount	Total Object Count	Total Object Count

Table 1-88. Historical XDB Metrics for Persistence (Continued)

Metric Key	Metric Name	Description
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
HisXDB|LockOperationErrorCount	Lock Operation Error Count	Lock Operation Error Count
HisXDB|DBCorruptionErrorCount	DB Corruption Error Count	DB Corruption Error Count
HisXDB|DBMaxSessionExceededCount	DB Maximum Sessions Exceeded Count	DB Maximum Sessions Exceeded Count
HisXDB|NumberWaitingForSession	Number of operations waiting for a session	Number of operations waiting for a session from the session pool
HisXDB|AvgWaitForSessionDuration	Average acquisition time from session pool	Average acquisition time from session pool
HisXDB|MinWaitForSessionDuration	Minimum acquisition time from session pool	Minimum acquisition time from session pool
HisXDB|MaxWaitForSessionDuration	Maximum acquisition time from session pool	Maximum acquisition time from session pool
HisXDB|TotalGetSessionCount	Total requests for a session from the session pool	Total requests for a session from the session pool
HisXDB|HisActivitySubmissionCount	HIS activity submission count	Number of Historical Inventory Service activities submitted
HisXDB|HisActivityCompletionCount	HIS activity completion count	Number of Historical Inventory Service activities completed
HisXDB|HisActivityCompletionDelayAvg	HIS activity average completion delay	The average amount of time from activity submission to completion
HisXDB|HisActivityCompletionDelayMax	HIS activity maximum completion delay	The maximum amount of time from activity submission to completion
HisXDB|HisActivityAbortedCount	HIS activity abort count	Number of Historical Inventory Service activities aborted

Remote Collector Metrics

vRealize Operations Manager collects metrics for the vRealize Operations Manager remote collector node objects.

Table 1-89. Remote Collector Metrics

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

Table 1-90. 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)

Table 1-90. Memory Metrics for the Remote Collector (Continued)

Metric Key	Metric Name	Description
mem|used	Used	Used
mem|total	Total	Total
mem|demand_gb	Estimated memory demand	Estimated memory demand

Table 1-91. Swap Metrics for the Remote Collector

Metric Key	Metric Name	Description
swap|total	Total	Total
swap|free	Free	Free
swap|used	Used	Used
swap|pageIn	Page in	Page in
swap|pageOut	Page out	Page out

Table 1-92. Resource limit Metrics for the Remote Collector

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

Table 1-93. Network Metrics for the Remote Collector

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

Table 1-93. Network Metrics for the Remote Collector (Continued)

Metric Key	Metric Name	Description
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 1-94. Network Interface Metrics for the Remote Collector

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

Table 1-95. Disk Filesystem Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|fileSystem|total	Total	Total
disk|fileSystem|available	Available	Available
disk|fileSystem|used	Used	Used
disk|fileSystem|files	Total file nodes	Total number of file nodes
disk|fileSystem|filesFree	Total free file nodes	Total free file nodes
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 1-96. Disk Installation Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|installation|used	Used	Used
disk|installation|total	Total	Total
disk|installation|available	Available	Available

Table 1-97. Disk Database Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|db|used	Used	Used
disk|db|total	Total	Total
disk|db|available	Available	Available

Table 1-98. Disk Log Metrics for the Remote Collector

Metric Key	Metric Name	Description
disk|log|used	Used	Used
disk|log|total	Total	Total
disk|log|available	Available	Available

Table 1-99. CPU Metrics for the Remote Collector

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

Table 1-100. Device Metrics for the Remote Collector

Metric Key	Metric Name	Description
device|iops	Reads/writes per second	Average number of read/write commands issued per second during the collection interval
device|await	Average transaction time	Average transaction time (milliseconds)

Table 1-101. Service Metrics for the Remote Collector

Metric Key	Metric Name	Description
service|proc|fdUsage	Total number of open file descriptors	Total number of open file descriptors (Linux). Total number of open handles (Windows)

Table 1-102. NTP Metrics for the Remote Collector

Metric Key	Metric Name	Description
ntp|serverCount	Configured server count	Configured server count
ntp|unreachableCount	Unreachable server count	Unreachable server count
ntp|unreachable	Unreachable	Is the NTP server unreachable. Value of 0 is reachable, 1 means the server was not reached or didn't respond.

Metrics for vRealize Automation

The vRealize Automation solution collects metrics about the total number of virtual machines deployed.

Table 1-103. Metrics

Metric Name	Description
Total Deployed	The total number of virtual machines deployed from the blueprint. Key: DeploymentCount | TotalDeployed

Managed Resources Object as a Filter in vRealize Automation

The vRealize Automation solution uses filters to display the VMware vCenter adapter objects that are managed by or have some association with vRealize Automation. Some of the dashboards have widgets that are configured to display only VMware vCenter adapter objects that vRealize Automation manages or is associated with. vRealize Automation uses an object called Managed Resources as a filter to display only those objects. All of these resources are placed under the Managed Resources object of type vRealize Automation Entity Status. The absence of this filter causes all VMware vCenter adapter objects to be displayed in the widgets. If you delete the Managed Resources object, the adapter re-creates the object, but the dashboards display incorrect information in the widgets that use this filter. If you delete the Managed Resources object, you must manually configure the widgets in the dashboard and select the Managed Resources object as a filter in each section that displays the VMware vCenter adapter objects.

Metrics for vSAN

vRealize Operations Manager collects metrics for vSAN objects.

In the menu, click **Environment > All Objects > vSAN Adapter**. Select one of the vSAN adapter objects listed and click the **All Metrics** tab.

Disk I/O and Disk Space Metrics for vSAN Disk Groups

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN disk groups.

Disk I/O metrics for the vSAN disk groups include:

- Disk I/O|Reads Per Second (IOPS)
- Disk I/O|Writes Per Second (IOPS
- Disk I/O|Max Observed Reads Per Second (IOPS)
- Disk I/O|Max Observed Writes Per Second (IOPS)
- Disk I/O|Throughput Read (bps)
- Disk I/O|Throughput Write (bps)
- Disk I/O|Average Read Latency (ms)
- Disk I/O|Average Write Latency (ms)
- Disk I/O|Read Count
- Disk I/O|Write Count
- Disk I/O|Average Device Latency
- Disk I/O|Average Device Read Latency
- Disk I/O|Average Device Write Latency
- Disk I/O|Total Bus Resets
- Disk I/O|Total Commands Aborted per second
- Disk I/O|Total Number of Errors

Disk space metrics for vSAN disk groups include:

- Disk Space|Capacity (bytes)
- Disk Space|Used (bytes)
- Disk Space|Usage (%)

Read Cache Metrics for vSAN Disk Groups

The vRealize Operations Manager collects metrics and performs capacity trend analysis on a hybrid vSAN read cache. Read Cache metrics are not collected for a vSAN all-flash configuration.

Read cache metrics for the vSAN disk group include:

- Read Cache|Hit Rate (%)
- Read Cache|Miss Rate Ratio
- Read Cache|Read Cache Reads Per Second (IOPS)
- Read Cache|Read Cache Read Latency (ms)
- Read Cache|Read Cache Read I/O Count
- Read Cache|Read Cache Writes Per Second (IOPS)
- Read Cache|Read Cache Write Latency (ms)
- Read Cache|Read Cache Write I/O Count

Write Buffer Metrics for vSAN Disk Groups

The vRealize Operations Manager collects the metrics you use to monitor the write buffer capacity of your vSAN disk groups.

A reasonably balanced system consumes a significant amount of write buffer. Before placing additional workload on the vSAN, check the write buffer metrics for the vSAN disk group.

- Write Buffer|Capacity (bytes)
- Write Buffer|Free (%)
- Write Buffer|Usage (%)
- Write Buffer|Used (byte)
- Write Buffer|Write Buffer Reads Per Second (IOPS)
- Write Buffer|Write Buffer Read Latency (ms)
- Write Buffer|Write Buffer Read I/O Count
- Write Buffer|Write Buffer Writes Per Second (IOPS)
- Write Buffer|Write Buffer Write Latency (ms)
- Write Buffer|Write Buffer Write I/O Count

Congestion Metrics for vSAN Disk Groups

The vRealize Operations Manager collects congestion metrics for the vSAN disk group.

- Congestion| Memory Congestion - Favorite
- Congestion| SSD Congestion - Favorite
- Congestion| IOPS Congestion - Favorite
- Congestion| Slab Congestion
- Congestion| Log Congestion
- Congestion| Comp Congestion

Metrics for vSAN Cluster

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN cluster.

Metrics for vSAN cluster include:

Component	Metrics
Component Limit	■ vSAN|Component Limit|Component Limit Used (%) ■ vSAN|Component Limit|Total Component Limit ■ vSAN|Component Limit|Used Component Limit
Disk Space	■ vSAN|Disk Space|Disk Space Used (%) ■ vSAN|Disk Space|Total Disk Space (GB) ■ vSAN|Disk Space|Used Disk Space (GB)
Read Cache	■ vSAN|Read Cache|Read Cache Reserved (%) ■ vSAN|Read Cache|Reserved Read Cache Size (GB) ■ vSAN|Read Cache|Total Read Cache Size (GB)

Component	Metrics
Performance	■ vSAN|Read Cache|Reads Per Second (IOPS) ■ vSAN|Read Cache|Read Throughput (KBps) ■ vSAN|Read Cache|Average Read Latency (ms) ■ vSAN|Read Cache|Writes Per Second (IOPS) ■ vSAN|Read Cache|Write Throughput (KBps) ■ vSAN|Read Cache|Average Write Latency (ms) ■ vSAN|Read Cache|Congestion ■ vSAN|Read Cache|Outstanding I/O ■ vSAN|Read Cache|Total IOPS ■ vSAN|Read Cache|Total Latency (ms) ■ vSAN|Read Cache|Total Throughput (KBps)
Deduplication And Compression Overview	■ vSAN|Deduplication And Compression Overview|Used Before ■ vSAN|Deduplication And Compression Overview|Used After ■ vSAN|Deduplication And Compression Overview|Savings ■ vSAN|Deduplication And Compression Overview|Ratio
Summary	■ Summary|Number of Cache Disks ■ Summary|Total Number of Capacity Disks ■ Summary|CPU Workload ■ Summary|Memory Workload ■ Summary|Total Number of Disk Groups ■ Summary|Total Active Alerts Count ■ Summary|Total Number of VMs ■ Summary|Total Number of Hosts ■ Summary|vSAN Cluster Capacity Remaining (%) ■ Summary|VSAN Cluster Storage Time Remaining ■ Summary|vSAN Capacity Disk Used

Metrics for vSAN Enabled Host

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN enabled host.

Metrics for vSAN enabled host include:

Component	Metrics
Component Limit	■ vSAN|Component Limit|Component Limit Used (%) ■ vSAN|Component Limit|Total Component Limit ■ vSAN|Component Limit|Used Component Limit
Disk Space	■ vSAN|Disk Space|Disk Space Used (%) ■ vSAN|Disk Space|Total Disk Space (GB) ■ vSAN|Disk Space|Used Disk Space (GB)
Read Cache	■ vSAN|Read Cache|Read Cache Reserved (%) ■ vSAN|Read Cache|Reserved Read Cache Size (GB) ■ vSAN|Read Cache|Total Read Cache Size (GB)

Metrics for vSAN Datastore

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN datastore.

Datastore I/O metrics for vSAN datastore include:

- Datastore I/O|Reads Per Second (IOPS)
- Datastore I/O|Read Rate (KBps)
- Datastore I/O|Read Latency (ms)
- Datastore I/O|Writes Per Second (IOPS)
- Datastore I/O|Write Rate (KBps)
- Datastore I/O|Write Latency (ms)
- Datastore I/O|Outstanding I/O requests
- Datastore I/O|Congestion

Metrics for vSAN Cache Disk

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN cache disk.

Metrics for vSAN cache disk include:

Component	Metrics
Performance	■ Performance|Bus Resets ■ Performance|Commands Aborted Per Second ■ Performance|Device Latency (ms) ■ Performance|Device Read Latency (ms) ■ Performance|Device Write Latency (ms) ■ Performance|Read Requests Per Second ■ Performance|Average Reads Per Second ■ Performance|Write Requests Per Second ■ Performance|Average Writes Per Second ■ Performance|Read Rate ■ Performance|Write Rate ■ Performance|Usage ■ Performance|HDD Errors
SCSI SMART Statistics	■ SCSI SMART Statistics|Health Status ■ SCSI SMART Statistics|Media Wearout Indicator ■ SCSI SMART Statistics|Write Error Count ■ SCSI SMART Statistics|Read Error Count ■ SCSI SMART Statistics|Power on Hours ■ SCSI SMART Statistics|Reallocated Sector Count ■ SCSI SMART Statistics|Raw Read Error Rate ■ SCSI SMART Statistics|Drive Temperature ■ SCSI SMART Statistics|Maximum Observed Drive Temperature ■ SCSI SMART Statistics|Drive Rated Max Temperature ■ SCSI SMART Statistics|Write Sectors TOT Count ■ SCSI SMART Statistics|Read Sectors TOT Count ■ SCSI SMART Statistics|Initial Bad Block Count

Component	Metrics
Capacity	■ vSAN|Health|Capacity|Total Disk Capacity (GB) ■ vSAN|Health|Capacity|Used Disk Capacity (GB)
Congestion Health	■ vSAN|Health|Congestion Health|Congestion Value
Performance	■ vSAN|Performance|Physical Layer Reads Per Second ■ vSAN|Performance|Physical Layer Writes Per Second ■ vSAN|Performance|Physical Layer Read Throughput (KBps) ■ vSAN|Performance|Physical Layer Write Throughput (KBps) ■ vSAN|Performance|Physical Layer Read Latency (ms) ■ vSAN|Performance|Physical Layer Write Latency (ms) ■ vSAN|Performance|Physical Layer Read Count ■ vSAN|Performance|Physical Layer Write Count ■ vSAN|Performance|Device Average Latency (ms) ■ vSAN|Performance|Guest Average Latency (ms)

Metrics for vSAN Capacity Disk

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN capacity disk.

Metrics for vSAN capacity disk include:

Component	Metrics
Performance	■ Performance|Bus Resets ■ Performance|Commands Aborted Per Second ■ Performance|Device Latency (ms) ■ Performance|Device Read Latency (ms) ■ Performance|Device Write Latency (ms) ■ Performance|Read Requests Per Second ■ Performance|Average Reads Per Second ■ Performance|Write Requests Per Second ■ Performance|Average Writes Per Second ■ Performance|Read Rate ■ Performance|Write Rate ■ Performance|Usage ■ Performance|HDD Errors
SCSI SMART Statistics	■ SCSI SMART Statistics|Health Status ■ SCSI SMART Statistics|Media Wearout Indicator ■ SCSI SMART Statistics|Write Error Count ■ SCSI SMART Statistics|Read Error Count ■ SCSI SMART Statistics|Power on Hours ■ SCSI SMART Statistics|Reallocated Sector Count ■ SCSI SMART Statistics|Raw Read Error Rate ■ SCSI SMART Statistics|Drive Temperature ■ SCSI SMART Statistics|Maximum Observed Drive Temperature ■ SCSI SMART Statistics|Drive Rated Max Temperature ■ SCSI SMART Statistics|Write Sectors TOT Count ■ SCSI SMART Statistics|Read Sectors TOT Count ■ SCSI SMART Statistics|Initial Bad Block Count
Capacity	■ vSAN|Health|Total Disk Capacity (GB) ■ vSAN|Health|Used Disk Capacity (GB)

Component	Metrics
Congestion Health	vSAN|Health|Congestion Value
Performance	■ vSAN|Performance|Physical Layer Reads Per Second ■ vSAN|Performance|Physical Layer Writes Per Second ■ vSAN|Performance|Physical Layer Read Throughput (KBps) ■ vSAN|Performance|Physical Layer Write Throughput (KBps) ■ vSAN|Performance|Physical Layer Read Latency (ms) ■ vSAN|Performance|Physical Layer Write Latency (ms) ■ vSAN|Performance|Physical Layer Read Count ■ vSAN|Performance|Physical Layer Write Count ■ vSAN|Performance|Device Average Latency (ms) ■ vSAN|Performance|Guest Average Latency (ms) ■ vSAN|Performance|vSAN Layer Reads Per Second ■ vSAN|Performance|vSAN Layer Writes Per Second ■ vSAN|Performance|vSAN Layer Read Latency (ms) ■ vSAN|Performance|vSAN Layer Write Latency (ms) ■ vSAN|Performance|vSAN Layer Read Count ■ vSAN|Performance|vSAN Layer Write Count

Properties for vSAN capacity disk include:

- Name
- Size
- Vendor
- Type
- Queue Depth

Metrics for vSAN World

The vRealize Operations Manager collects the metrics you use to monitor the performance of your vSAN world.

Metrics for vSAN world include:

- Summary|Total Number of VMs
- Summary|Total Number of Hosts
- Summary|Total IOPS
- Summary|Total Latency
- Summary|Total Number of Clusters
- Summary|Total Number of DiskGroups
- Summary|Total Number of Cache Disks
- Summary|Total Number of Capacity Disks
- Summary|Total Number of Datastores
- Summary|Total vSAN Disk Capacity (TB)
- Summary|Total vSAN Disk Capacity Used (TB)
- Summary|Remaining Capacity (TB)
- Summary|Remaining Capacity (%)

- Summary | Total Savings by Deduplication and Compression (GB)

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

vRealize Operations Manager collects metrics for the object types in the Operating Systems and Remote Service Monitoring plug-ins.

Due to rounding in metric time calculation, there can be situations in which the Resource Availability metric is rounded up. Rounding up the metric appears as gaps in the metrics reported by the End Point Operations Management agent. However, the metrics are fully reported.

Operating Systems Plug-in Metrics

The Operating Systems plug-in collects metrics for object types such Linux, AIX, Solaris, and Windows. The Operating Systems plug-in also collects metrics for Windows services, Script services, and Multiprocess services.

AIX Metrics

The Operating Systems Plug-in discovers the metrics for the AIX object type. AIX 6.1 and 7.1 are supported.

Table 1-104. AIX metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	True
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False

Table 1-104. AIX metrics (Continued)

Name	Category	KPI
Cpu Wait Time	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False
Nfs Server V3 Read per Minute	UTILIZATION	False
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Symlink	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False

Table 1-104. AIX metrics (Continued)

Name	Category	KPI
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False
Nfs Server V3 Readdir per Minute	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Page Major faults	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Page Faults	UTILIZATION	False
Percent Used Swap	UTILIZATION	True
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Pages In	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False

Table 1-104. AIX metrics (Continued)

Name	Category	KPI
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Linux Metrics

The Operating Systems Plug-in discovers the metrics for the Linux object type.

Table 1-105. Linux Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False

Table 1-105. Linux Metrics (Continued)

Name	Category	KPI
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Stolen	UTILIZATION	False
Cpu Wait Time	UTILIZATION	False
Cpu Irq Time per Minute	UTILIZATION	False
Cpu SoftIrq Time	UTILIZATION	False
Cpu Stolen Time per Minute	UTILIZATION	False
Cpu Stolen Time	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Irq	UTILIZATION	False
Cpu SoftIrq Time per Minute	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Irq Time	UTILIZATION	False
Cpu SoftIrq	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Free Memory (+ buffers/cache)	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False

Table 1-105. Linux Metrics (Continued)

Name	Category	KPI
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False
Nfs Server V3 Read per Minute	UTILIZATION	False
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False

Table 1-105. Linux Metrics (Continued)

Name	Category	KPI
Nfs Server V3 Readdir per Minute	UTILIZATION	False
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Page Major faults	UTILIZATION	False
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Page Faults	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
Used Memory (- buffers/cache)	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Solaris Metrics

The Operating Systems Plug-in discovers the metrics for the Solaris object type. Solaris x86 and SPARC are supported.

Table 1-106. Solaris Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False

Table 1-106. Solaris Metrics (Continued)

Name	Category	KPI
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Wait Time	UTILIZATION	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Wait Time per Minute	UTILIZATION	False
Cpu Idle	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Wait	UTILIZATION	False
Cpu Nice	UTILIZATION	False
Free Memory	UTILIZATION	False
Load Average 15 Minutes	UTILIZATION	False
Load Average 5 Minutes	UTILIZATION	False
Load Average 1 Minute	UTILIZATION	False
Nfs Server V3 Readlink per Minute	UTILIZATION	False
Nfs Server V3 Readdirplus per Minute	UTILIZATION	False
Nfs Server V3 Commit per Minute	UTILIZATION	False
Nfs Server V3 Access	UTILIZATION	False
Nfs Server V3 Access per Minute	UTILIZATION	False
Nfs Server V3 Remove	UTILIZATION	False
Nfs Server V3 Rename per Minute	UTILIZATION	False
Nfs Server V3 Fsstat per Minute	UTILIZATION	False
Nfs Server V3 Create per Minute	UTILIZATION	False
Nfs Server V3 Mkdir per Minute	UTILIZATION	False
Nfs Server V3 Mknod	UTILIZATION	False
Nfs Server V3 Read per Minute	UTILIZATION	False
Nfs Server V3 Fsstat	UTILIZATION	False
Nfs Server V3 Link	UTILIZATION	False
Nfs Server V3 Write	UTILIZATION	False
Nfs Server V3 Remove per Minute	UTILIZATION	False
Nfs Server V3 Lookup per Minute	UTILIZATION	False
Nfs Server V3 Link per Minute	UTILIZATION	False
Nfs Server V3 Rmdir per Minute	UTILIZATION	False
Nfs Server V3 Mkdir	UTILIZATION	False
Nfs Server V3 Mknod per Minute	UTILIZATION	False
Nfs Server V3 Getattr per Minute	UTILIZATION	False
Nfs Server V3 Null	UTILIZATION	False
Nfs Server V3 Readdirplus	UTILIZATION	False
Nfs Server V3 Lookup	UTILIZATION	False
Nfs Server V3 Pathconf	UTILIZATION	False

Table 1-106. Solaris Metrics (Continued)

Name	Category	KPI
Nfs Server V3 Readlink	UTILIZATION	False
Nfs Server V3 Write per Minute	UTILIZATION	False
Nfs Server V3 Readdir	UTILIZATION	False
Nfs Server V3 Setattr per Minute	UTILIZATION	False
Nfs Server V3 Setattr	UTILIZATION	False
Nfs Server V3 Read	UTILIZATION	False
Nfs Server V3 Pathconf per Minute	UTILIZATION	False
Nfs Server V3 Symlink per Minute	UTILIZATION	False
Nfs Server V3 Symlink	UTILIZATION	False
Nfs Server V3 Fsinfo per Minute	UTILIZATION	False
Nfs Server V3 Fsinfo	UTILIZATION	False
Nfs Server V3 Getattr	UTILIZATION	False
Nfs Server V3 Rmdir	UTILIZATION	False
Nfs Server V3 Readdir per Minute	UTILIZATION	False
Nfs Server V3 Create	UTILIZATION	False
Nfs Server V3 Rename	UTILIZATION	False
Nfs Server V3 Commit	UTILIZATION	False
Nfs Server V3 Null per Minute	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Page Major faults	UTILIZATION	False
Page Major faults per Second	UTILIZATION	False
Page Faults per Second	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Page Faults	UTILIZATION	False
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False

Table 1-106. Solaris Metrics (Continued)

Name	Category	KPI
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	False
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Microsoft Windows Metrics

The Operating Systems Plug-in discovers the metrics for the Microsoft Windows object type. Microsoft Windows Server 2012 R2 and 2008 R2 are supported.

Table 1-107. Microsoft Windows Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
System Uptime	AVAILABILITY	False
Avg. Disk sec/Transfer	THROUGHPUT	False
File System Reads/Writes	THROUGHPUT	False
File System Reads/Writes per Minute	THROUGHPUT	False
Tcp Attempt Fails	THROUGHPUT	False
Tcp State Established	THROUGHPUT	False
Tcp Estab Resets per Minute	THROUGHPUT	False
Tcp Retrans Segs	THROUGHPUT	False
Tcp State LISTEN	THROUGHPUT	False
Tcp State CLOSING	THROUGHPUT	False
Tcp State SYN_SENT	THROUGHPUT	False
Tcp State TIME_WAIT	THROUGHPUT	False
Tcp State SYN_RECV	THROUGHPUT	False
Tcp In Errs per Minute	THROUGHPUT	False
Tcp Out Segs per Minute	THROUGHPUT	False
Tcp Passive Opens per Minute	THROUGHPUT	False
Tcp Out Segs	THROUGHPUT	False
Tcp Estab Resets	THROUGHPUT	False
Tcp Active Opens	THROUGHPUT	False
Tcp Outbound Connections	THROUGHPUT	False
Tcp Curr Estab	THROUGHPUT	False
Tcp In Errs	THROUGHPUT	False

Table 1-107. Microsoft Windows Metrics (Continued)

Name	Category	KPI
Tcp Inbound Connections	THROUGHPUT	False
Tcp Active Opens per Minute	THROUGHPUT	False
Tcp Out Rsts per Minute	THROUGHPUT	False
Tcp In Segs	THROUGHPUT	False
Tcp Retrans Segs per Minute	THROUGHPUT	False
Tcp Passive Opens	THROUGHPUT	False
Tcp Out Rsts	THROUGHPUT	False
Tcp State FIN_WAIT1	THROUGHPUT	False
Tcp State FIN_WAIT2	THROUGHPUT	False
Tcp State CLOSE_WAIT	THROUGHPUT	False
Tcp In Segs per Minute	THROUGHPUT	False
Tcp State CLOSE	THROUGHPUT	False
Tcp State LAST_ACK	THROUGHPUT	False
Tcp Attempt Fails per Minute	THROUGHPUT	False
Cpu Idle Time	UTILIZATION	False
Cpu Idle Time per Minute	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Free Memory	UTILIZATION	False
Memory Page Faults/sec	UTILIZATION	False
Memory System Driver Resident Bytes	UTILIZATION	False
Memory Available Bytes	UTILIZATION	False
Memory System Driver Total Bytes	UTILIZATION	False
Memory % Committed Bytes In Use	UTILIZATION	False
Memory Standby Cache Core Bytes	UTILIZATION	False
Memory Transition Pages RePurposed/sec	UTILIZATION	False
Memory Write Copies/sec	UTILIZATION	False
Memory Available KBytes	UTILIZATION	False
Memory Page Reads/sec	UTILIZATION	False
Memory Committed Bytes	UTILIZATION	False
Memory Pool Nonpaged Bytes	UTILIZATION	False
Memory System Code Resident Bytes	UTILIZATION	False
Memory Page Writes/sec	UTILIZATION	False
Memory Available MBytes	UTILIZATION	False
Memory Standby Cache Normal Priority Bytes	UTILIZATION	False
Memory Pages/sec	UTILIZATION	False
Memory Modified Page List Bytes	UTILIZATION	False
Memory Cache Faults/sec	UTILIZATION	False
Memory Pool Nonpaged Allocs	UTILIZATION	False
Memory System Code Total Bytes	UTILIZATION	False

Table 1-107. Microsoft Windows Metrics (Continued)

Name	Category	KPI
Memory Pool Paged Allocs	UTILIZATION	False
Memory Pages Input/sec	UTILIZATION	False
Memory Pool Paged Bytes	UTILIZATION	False
Memory Pool Paged Resident Bytes	UTILIZATION	False
Memory Cache Bytes	UTILIZATION	False
Memory Standby Cache Reserve Bytes	UTILIZATION	False
MemoryFreeSystemPageTableEntries	UTILIZATION	False
Memory Free %26 Zero Page List Bytes	UTILIZATION	False
Memory System Cache Resident Bytes	UTILIZATION	False
Memory Cache Bytes Peak	UTILIZATION	False
Memory Commit Limit	UTILIZATION	False
Memory Transition Faults/sec	UTILIZATION	False
Memory Pages Output/sec	UTILIZATION	False
Number of CPUs	UTILIZATION	False
Percent Free Swap	UTILIZATION	False
Percent Free Memory	UTILIZATION	False
Percent Used Memory	UTILIZATION	True
Percent Used Swap	UTILIZATION	True
Running Processes	UTILIZATION	False
Sleeping Processes	UTILIZATION	False
Stopped Processes	UTILIZATION	False
Swap Pages Out per Minute	UTILIZATION	False
Swap Pages In per Minute	UTILIZATION	False
Swap Free	UTILIZATION	False
Swap Pages Out	UTILIZATION	False
Swap Used	UTILIZATION	False
Swap Total	UTILIZATION	False
Swap Pages In	UTILIZATION	False
System Cpu	UTILIZATION	False
System Cpu Time per Minute	UTILIZATION	False
System Cpu Time	UTILIZATION	False
Total disk capacity	UTILIZATION	False
Total Processes	UTILIZATION	False
Total Memory	UTILIZATION	True
Total disk usage	UTILIZATION	False
User Cpu Time	UTILIZATION	False
User Cpu	UTILIZATION	False
User Cpu Time per Minute	UTILIZATION	False

Table 1-107. Microsoft Windows Metrics (Continued)

Name	Category	KPI
Used Memory	UTILIZATION	False
Zombie Processes	UTILIZATION	False

Windows Service Metrics

The Operating Systems Plug-in discovers the metrics for Windows service.

Table 1-108. Windows Services Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Start Time	AVAILABILITY	False
Start Type	AVAILABILITY	False
Cpu User Time	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Total Time per Minute	UTILIZATION	False
Cpu System Time per Minute	UTILIZATION	False
Cpu Total Time	UTILIZATION	False
Cpu User Time per Minute	UTILIZATION	False
Cpu System Time	UTILIZATION	False
Memory Size	UTILIZATION	True
Open Handles	UTILIZATION	False
Resident Memory Size	UTILIZATION	False
Threads	UTILIZATION	False

If you stop an End Point Operations Management agent by using Windows Services, and remove the `data` directory from inside the agent installation directory, when you start the agent again, using Windows Services, no metrics are collected. If you are deleting the `data` directory, do not use Windows Services to stop and start an End Point Operations Management agent. Stop the agent using `epops-agent.bat stop`. Delete the `data` directory, then start the agent using `epops-agent.bat start`.

Script Metrics

The Operating Systems Plug-in discovers the metrics for the Script service.

Table 1-109. Script Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Execution Time	THROUGHPUT	True
Result Value	UTILIZATION	True

Multiprocess Service Metrics

The Operating Systems Plug-in discovers the metrics for the Multiprocess service.

Table 1-110. Multiprocess Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Cpu User Time	UTILIZATION	False
Cpu Usage	UTILIZATION	True
Cpu Total Time per Minute	UTILIZATION	False
Cpu System Time per Minute	UTILIZATION	False
Cpu Total Time	UTILIZATION	False
Cpu User Time per Minute	UTILIZATION	False
Cpu System Time	UTILIZATION	False
Memory Size	UTILIZATION	True
Number of Processes	UTILIZATION	False
Resident Memory Size	UTILIZATION	False

NFS Metrics

The End Point Operations Management agents collect metrics for the NFS-mounted file systems.

The following metrics are collected.

Name	Category
Resource Availability	Availability
Use Percent (%)	Utilization
Total Bytes Free (KB)	Utilization

Remote Service Monitoring Plug-in Metrics

The Remote Service Monitoring plug-in collects metrics for object types such HTTP Check, TCP Check, and ICMP Check.

HTTP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the HTTP Check object type.

Table 1-111. HTTP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Last Modified	AVAILABILITY	False
State CLOSE	THROUGHPUT	False
State CLOSE_WAIT	THROUGHPUT	False
State ESTABLISHED	THROUGHPUT	False
Inbound Connections	THROUGHPUT	False
State TIME_WAIT	THROUGHPUT	False

Table 1-111. HTTP Check Metrics (Continued)

Name	Category	KPI
All Inbound Connections	THROUGHPUT	False
State SYN_SENT	THROUGHPUT	False
State FIN_WAIT2	THROUGHPUT	False
Outbound Connections	THROUGHPUT	False
State LAST_ACK	THROUGHPUT	False
Response Time	THROUGHPUT	True
State CLOSING	THROUGHPUT	False
All Outbound Connections	THROUGHPUT	False
State SYN_RECV	THROUGHPUT	False
State FIN_WAIT1	THROUGHPUT	False
Response Code	UTILIZATION	True

ICMP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the ICMP Check object type.

Table 1-112. ICMP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Response Time	THROUGHPUT	True

TCP Check Metrics

The Remote Service Monitoring Plug-in discovers the metrics for the TCP Check object type.

Table 1-113. TCP Check Metrics

Name	Category	KPI
Resource Availability	AVAILABILITY	True
Response Time	THROUGHPUT	True
State CLOSE	THROUGHPUT	False
State CLOSE_WAIT	THROUGHPUT	False
State ESTABLISHED	THROUGHPUT	False
Inbound Connections	THROUGHPUT	False
State TIME_WAIT	THROUGHPUT	False
All Inbound Connections	THROUGHPUT	False
State SYN_SENT	THROUGHPUT	False
State FIN_WAIT2	THROUGHPUT	False
Outbound Connections	THROUGHPUT	False
State LAST_ACK	THROUGHPUT	False
State CLOSING	THROUGHPUT	False
All Outbound Connections	THROUGHPUT	False

Table 1-113. TCP Check Metrics (Continued)

Name	Category	KPI
State SYN_RECV	THROUGHPUT	False
State FIN_WAIT1	THROUGHPUT	False

Property Definitions in vRealize Operations Manager 2

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

vRealize Operations Manager uses adapters to collect properties for target objects in your environment. Property definitions for all objects connected through the vCenter adapter are provided. The properties collected depend on the objects in your environment.

You can add symptoms based on properties to an alert definition so that you are notified if a change occurs to properties on your monitored objects. For example, disk space is a hardware property of a virtual machine. You can use disk space to define a symptom that warns you when the value falls below a certain numeric value. See the *vRealize Operations Manager User Guide*.

vRealize Operations Manager generates Object Type Classification and Subclassification properties for every object. You can use object type classification properties to identify whether an object is an adapter instance, custom group, application, tier, or a general object with property values *ADAPTER_INSTANCE*, *GROUP*, *BUSINESS_SERVICE*, *TIER*, or *GENERAL*, respectively.

This chapter includes the following topics:

Properties for vCenter Server Components

The VMware vSphere solution is installed with vRealize Operations Managerand includes the vCenter adapter. vRealize Operations Manager uses the vCenter adapter to collect properties for objects in the vCenter Server system.

vCenter Server components are listed in the `describe.xml` file for the vCenter adapter. The following example shows the runtime property `memoryCap` or Memory Capacity for the virtual machine in the `describe.xml`.

```
<ResourceGroup instanced="false" key="runtime" nameKey="5300" validation="">
   <ResourceAttribute key="memoryCap" nameKey="1780" dashboardOrder="200" dataType="float"
                      defaultMonitored="true" isDiscrete="false" isRate="false" maxVal=""
                      minVal="" isProperty="true" unit="kb"/>
</ResourceGroup>
```

The `ResourceAttribute` element includes the name of the property that appears in the UI and is documented as a Property Key. `isProperty = "true"` indicates that `ResourceAttribute` is a property.

vCenter Server Properties

vRealize Operations Manager collects summary and event properties for vCenter Server system objects.

Table 2-1. Summary Properties Collected for vCenter Server System 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 Server System 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 Server System Objects

Property Key	Property Name	Description
CustomFieldManager|CustomFieldDef	Custom Field Def	Custom Field Def for VCenter Tagging information at Adapter level.

Virtual Machine Properties

vRealize Operations Manager collects configuration, runtime, CPU, memory, network I/O, summary, guest file system, and properties about datastore use for virtual machine objects.

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

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

Table 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. Configuration Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
config|name	Name	Name
config|guestFullName	Guest Fullname	Guest OS full name configured by the user.
config|hardware|numCpu	Number of virtual CPUs	Number of virtual CPUs
config|hardware|memoryKB	Memory	Memory

Table 2-6. Configuration Properties Collected for Virtual Machine Objects (Continued)

Property Key	Property Name	Description
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)
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)

Table 2-6. Configuration Properties Collected for Virtual Machine Objects (Continued)

Property Key	Property Name	Description
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_memsfss	Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable)	Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable)
config|security|disable_unexposed_features_protocolhandler	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)	Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable)
config|security|disable_unexposed_features_shellaction	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)	Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable)
config|security|disable_unexposed_features_toporequest	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)	Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable)
config|security|disable_unexposed_features_trashfolderstate	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)	Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable)
config|security|disable_unexposed_features_trayicon	Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable)	Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable)
config|security|disable_unexposed_features_unity	Disable unexposed features - unity (isolation.tools.unity.disable)	Disable unexposed features - unity (isolation.tools.unity.disable)
config|security|disable_unexposed_features_unity_interlock	Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable)	Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable)
config|security|disable_unexposed_features_unity_taskbar	Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable)	Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable)
config|security|disable_unexposed_features_unity_unityactive	Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable)	Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable)
config|security|disable_unexposed_features_unity_windowcontents	Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable)	Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable)

Table 2-6. Configuration Properties Collected for Virtual Machine Objects (Continued)

Property Key	Property Name	Description
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 (solation.tools.guestDnDVersionSet.disable)	Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable)
config|security|disable_vix_messages	Disable VIX messages from the VM (isolation.tools.vixMessage.disable)	Disable VIX messages from the VM (isolation.tools.vixMessage.disable)
config|security|enable_vga_only_mode	Disable all but VGA mode on virtual machines (svga.vgaOnly)	Disable all but VGA mode on virtual machines (svga.vgaOnly)
config|security|limit_console_connection	Limit number of console connections (RemoteDisplay.maxConnection)	Limit number of console connections (RemoteDisplay.maxConnection)
config|security|limit_log_number	Limit number of log files (log.keepOld)	Limit number of log files (log.keepOld)
config|security|limit_log_size	Limit log file size (log.rotateSize)	Limit log file size (log.rotateSize)
config|security|limit_setinfo_size	Limit VMX file size (tools.setInfo.sizeLimit)	Limit VMX file size (tools.setInfo.sizeLimit)
config|security|enable_console_VNC	Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled)	Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled)
config|security|disable_device_interaction_connect	Disable unauthorized removal, connection of devices (isolation.device.connectable.disable)	Disable unauthorized removal, connection of devices (isolation.device.connectable.disable)
config|security|disable_device_interaction_edit	Disable unauthorized modification of devices (isolation.device.edit.disable)	Disable unauthorized modification of devices (isolation.device.edit.disable)
config|security|enable_host_info	Enable send host information to guests (tools.guestlib.enableHostInfo)	Enable send host information to guests (tools.guestlib.enableHostInfo)
config|security|network_filter_enable	Enable dvfilter network APIs (ethernetX.filterY.name)	Enable dvfilter network APIs (ethernetX.filterY.name)
config|security|vmsafe_cpumem_agentaddress	VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress)	VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress)
config|security|vmsafe_cpumem_agentport	VMsafe CPU/memory APIs - port number (vmsafe.agentPort)	VMsafe CPU/memory APIs - port number (vmsafe.agentPort)
config|security|vmsafe_cpumem_enable	Enable VMsafe CPU/memory APIs (vmsafe.enable)	Enable VMsafe CPU/memory APIs (vmsafe.enable)
config|security|disconnect_devices_floppy	Disconnect floppy drive	Disconnect floppy drive
config|security|disconnect_devices_cd	Disconnect CD-ROM	Disconnect CD-ROM
config|security|disconnect_devices_usb	Disconnect USB controller	Disconnect USB controller

Table 2-6. Configuration Properties Collected for Virtual Machine Objects (Continued)

Property Key	Property Name	Description
config|security|disconnect_devices_parallel	Disconnect parallel port	Disconnect parallel port
config|security|disconnect_devices_serial	Disconnect serial port	Disconnect serial port

Note Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

For more information on the *vSphere Hardening Guide* alerts, see the *vRealize Operations Manager User Guide*.

Table 2-7. Runtime Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
runtime|memoryCap	Memory Capacity	Memory Capacity

Table 2-8. CPU Usage Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
cpu|limit	CPU limit	CPU limit
cpu|reservation	CPU reservation	CPU reservation
cpu|speed	CPU	CPU Speed
cpu|cpuModel	CPU Model	CPU Model

Table 2-9. Memory Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
mem|host_reservation	VM Reservation	Mem Machine Reservation
mem|host_limit	VM Limit	Mem Machine Limit

Table 2-10. Network Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
net|mac_address	Mac Address	Mac Address
net|ip_address	IP Address	IP Address
net|subnet_mask	Subnet Mask	Subnet Mask
net|default_gateway	Default Gateway	Default Gateway
net|nvp_vm_uuid	NVP VM UUID	NVP VM UUID

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

Property Key	Property Name	Description
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name
summary|parentCluster	Parent Cluster	Parent Cluster
summary|parentHost	Parent Host	Parent Host
summary|parentDatacenter	Parent Datacenter	Parent Datacenter
summary|parentVcenter	Parent Vcenter	Parent Vcenter

Table 2-11. Summary Properties Collected for Virtual Machine Objects (Continued)

Property Key	Property Name	Description
summary|guest|fullName	Guest OS Full Name	Guest OS Full Name as identified by VMware tools
summary|guest|ipAddress	Guest OS IP Address	Guest OS IP Address
summary|guest|toolsRunningStatus	Tools Running Status	Guest Tools Running Status
summary|guest|toolsVersionStatus2	Tools Version Status	Guest Tools Version Status 2
summary|guest|vrealize_operations_agent_id	vRealize Operations Agent ID	An ID to identify a VM in Agent Adapter's world
summary|guest|vrealize_operations_euc_agent_id	vRealize Operations Euc Agent ID	An ID to identify a VM in Agent Adapter's world
summary|config|numEthernetCards	Number of NICs	Number of NICs
summary|config|isTemplate	VM Template	Indicates whether it is a VM Template
summary|runtime|powerState	Power State	Power State
summary|runtime|connectionState	Connection State	Connection State

Table 2-12. Datastore Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
datastore|maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore|maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Table 2-13. Guest File System Properties Collected for Virtual Machine Objects

Property Key	Property Name	Description
guestfilesystem|capacity_property	Guest File System Capacity Property	Total capacity of guest file system as a property, reported for each file system.
guestfilesystem|capacity_property_total	Total Guest File System Capacity Property	Overall total capacity of guest file system as a property, reported across all file systems.

Host System Properties

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

Table 2-14. 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

Table 2-14. Configuration Properties Collected for Host System Objects (Continued)

Property Key	Property Name	Description
config|network|dnsserver	DNS Server	List of DNS Servers
config|product|productLineId	Product Line ID	Product Line ID
config|product|apiVersion	API Version	API Version
config|storageDevice|plugStoreTopology|numberofPath	Total number of Path	Total number of storage paths
config|storageDevice|multipathInfo|numberofActivePath	Total number of Active Path	Total number of active storage paths
config|storageDevice|multipathInfo|multipathPolicy	Multipath Policy	Multipath Policy
config|hyperThread|available	Available	Indicates whether hyperthreading is supported by the server
config|hyperThread|active	Active	Indicates whether hyperthreading is active
config|ntp|server	NTP Servers	NTP Servers
config|security|ntpServer	NTP server	NTP server
config|security|enable_ad_auth	Enable active directory authentication	Enable active directory authentication
config|security|enable_chap_auth	Enable mutual chap authentication	Enable mutual chap authentication
config|security|enable_auth_proxy	Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate)	Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate)
config|security|syslog_host	Remote log host (Syslog.global.logHost)	Remote log host (Syslog.global.logHost)
config|security|dcui_access	Users who can override lock down mode and access the DCUI (DCUI.Access)	Users who can override lock down mode and access the DCUI (DCUI.Access)
config|security|shell_interactive_timeout	Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut)	Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut)
config|security|shell_timeout	Shell timeout (UserVars.ESXiShellTimeOut)	Shell timeout (UserVars.ESXiShellTimeOut)
config|security|dvfilter_bind_address	Dvfilter bind ip address (Net.DVFilterBindIpAddress)	Dvfilter bind ip address (Net.DVFilterBindIpAddress)
config|security|syslog_dir	Log directory (Syslog.global.logDir)	Log directory (Syslog.global.logDir)
config|security|firewallRule|allowedHosts	Allowed hosts	Allowed hosts in the firewall configuration
config|security|service|isRunning	Running	Indicates whether a service is running or not. Services are: Direct Console UI, ESXi shell, SSH, or NTP Daemon.
config|security|service|ruleSet	Ruleset	Ruleset for each service.
config|security|service|policy	Policy	Policy for each service.

Note Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

For more information on the *vSphere Hardening Guide* alerts, see the *vRealize Operations Manager User Guide*.

Table 2-15. Hardware Properties Collected for Host System Objects

Property Key	Property Name	Description
hardware|memorySize	Memory Size	Memory Size
hardware|cpuInfo|numCpuCores	Number of CPU Cores	Number of CPU Cores
hardware|cpuInfo|hz	CPU Speed per Core	CPU Speed per Core
hardware|cpuInfo|numCpuPackages	Number of CPU Packages	Number of CPU Packages
hardware|cpuInfo|powerManagementPolicy	Active CPU Power Management Policy	Active CPU Power Management Policy
hardware|cpuInfo|powerManagementTechnology	Power Management Technology	Power Management Technology
hardware|cpuInfo|biosVersion	BIOS Version	BIOS Version

Table 2-16. 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 2-17. 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 2-18. 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 2-19. 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

Table 2-19. Network Properties Collected for Host System Objects (Continued)

Property Key	Property Name	Description
net|discoveryProtocol|cdp|portName	Port Name	Port Name
net|discoveryProtocol|cdp|vlan	VLAN	VLAN
net|discoveryProtocol|cdp|mtu	MTU	MTU
net|discoveryProtocol|cdp|hardwarePlatform	Hardware Platform	Hardware Platform
net|discoveryProtocol|cdp|softwareVersion	Software Version	Software Version
net|discoveryProtocol|cdp|timeToLive	Time to Live	Time to Live
net|discoveryProtocol|lldp|managementIpAddress	Management IP Address	Management IP Address
net|discoveryProtocol|lldp|systemName	System Name	System Name
net|discoveryProtocol|lldp|portName	Port Name	Port Name
net|discoveryProtocol|lldp|vlan	VLAN	VLAN
net|discoveryProtocol|lldp|timeToLive	Time to Live	Time to Live

Table 2-20. 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 2-21. 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 2-22. Datastore Properties Collected for Host System Objects

Property Key	Property Name	Description
datastore|maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore|maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests

Table 2-22. Datastore Properties Collected for Host System Objects (Continued)

Property Key	Property Name	Description
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Cluster Compute Resource Properties

vRealize Operations Manager collects configuration and summary properties for cluster compute resource objects.

Table 2-23. Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 2-24. Summary Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
summary|parentDatacenter	Parent Datacenter	Parent Datacenter
summary|parentVcenter	Parent Vcenter	Parent Vcenter
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Table 2-25. DR, DAS, and DPM Configuration Properties Collected for Cluster Compute Resource Objects

Property Key	Property Name	Description
configuration|drsconfig|enabled	Enabled	Indicates whether DRS is enabled
configuration|drsconfig|defaultVmBehavior	Default DRS Behaviour	Default DRS Behaviour
configuration|drsconfig|affinityRules	Affinity Rules	DRS Affinity Rules
configuration|dasconfig|enabled	HA Enabled	HA Enabled
configuration|dasconfig|admissionControlEnabled	Admission Control Enabled	Admission Control Enabled
configuration|dpmconfiginfo|enabled	DPM Enabled	DPM Enabled
configuration|dpmconfiginfo|defaultDpmBehavior	Default DPM Behaviour	Default DPM Behaviour

DRS properties are collected for disaster recovery. DAS properties are collected for high availability service, formerly distributed availability service. DPM properties are collected for distributed power management.

Resource Pool Properties

vRealize Operations Manager collects configuration, CPU, memory, and summary properties for resource pool objects.

Table 2-26. Configuration Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
config|name	Name	Name
config|cpuAllocation|reservation	Reservation	CPU reservation

Table 2-26. Configuration Properties Collected for Resource Pool Objects (Continued)

Property Key	Property Name	Description
config|cpuAllocation|limit	Limit	CPU limit
config|cpuAllocation|expandableReservation	Expandable Reservation	CPU expandable reservation
config|cpuAllocation|shares|shares	Shares	CPU shares
config|memoryAllocation|reservation	Reservation	Memory reservation
config|memoryAllocation|limit	Limit	Memory limit
config|memoryAllocation|expandableReservation	Expandable Reservation	Memory expandable reservation
config|memoryAllocation|shares|shares	Shares	Memory shares

Table 2-27. CPU Usage Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
cpu|limit	CPU Limit	CPU Limit
cpu|reservation	CPU reservation	CPU Reservation
cpu|expandable_reservation	CPU expandable reservation	CPU Expandable Reservation
cpu|shares	CPU Shares	CPU Shares
cpu|corecount_provisioned	Provisioned vCPU(s)	Provisioned vCPU(s)

Table 2-28. 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 2-29. Summary Properties Collected for Resource Pool Objects

Property Key	Property Name	Description
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Data Center Properties

vRealize Operations Manager collects configuration and summary properties for data center objects.

Table 2-30. Configuration Properties Collected for Data Center Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 2-31. Summary Properties Collected for Data Center Objects

Property Key	Property Name	Description
summary|parentVcenter	Parent Vcenter	Parent Vcenter
summary|customTag|customTagValue	Value	Custom Tag Value
summary|tag	vSphere Tag	vSphere Tag Name

Storage Pod Properties

vRealize Operations Manager collects configuration and summary properties for storage pod objects.

Table 2-32. Configuration Properties Collected for Storage Pod Objects

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

VMware Distributed Virtual Switch Properties

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

Table 2-33. Configuration Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 2-34. Capability Properties Collected for VMware Distributed Virtual Switch Objects

Property Key	Property Name	Description
capability|nicTeamingPolicy	NIC Teaming Policy	NIC Teaming Policy

Distributed Virtual Port Group Properties

vRealize Operations Manager collects configuration and summary properties for distributed virtual port group objects.

Table 2-35. Configuration Properties Collected for Distributed Virtual Port Group Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 2-36. Summary Properties Collected for Distributed Virtual Port Group Objects

Property Key	Property Name	Description
summary|active_uplink_ports	Active DV uplinks	Active DV uplinks

Datastore Properties

vRealize Operations Manager collects configuration, summary, and properties about datastore use for datastore objects.

Table 2-37. Configuration Properties Collected for Datastore Objects

Property Key	Property Name	Description
config|name	Name	Name

Table 2-38. Summary Properties Collected for Datastore Objects

Property Key	Property Name	Description
summary|diskCapacity	Disk Capacity	Disk Capacity
summary|isLocal	Is Local	Is local datastore
summary|customTag|customTagValue	Value	Custom Tag Value
summary|accessible	Datastore Accessible	Datastore Accessible

Table 2-39. Datastore Properties Collected for Datastore Objects

Property Key	Property Name	Description
datastore|hostcount	Host Count	Host Count
datastore|hostScsiDiskPartition	Host SCSI Disk Partition	Host SCSI Disk Partition
datastore|maxObservedNumberRead	Highest Observed Number of Read Requests	Highest Observed Number of Read Requests
datastore|maxObservedRead	Highest Observed Read Rate	Highest Observed Read Rate (KBps)
datastore|maxObservedReadLatency	Highest Observed Read Latency	Highest Observed Read Latency
datastore|maxObservedNumberWrite	Highest Observed Number of Write Requests	Highest Observed Number of Write Requests
datastore|maxObservedWrite	Highest Observed Write Rate	Highest Observed Write Rate (KBps)
datastore|maxObservedWriteLatency	Highest Observed Write Latency	Highest Observed Write Latency
datastore|maxObservedOIO	Highest Observed Outstanding Requests	Highest Observed Outstanding Requests

Self-Monitoring Properties for vRealize Operations Manager

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

Analytics Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager analytics service.

Table 2-40. Properties Collected for Analytics Service Objects

Property Key	Property Name	Description
HAEnabled	HA Enabled	Indicates HA is enabled with a value of 1, disabled with a value of 0.
ControllerDBRole	Role	Indicates persistence service role for the controller: 0 – Master, 1 – Replica, 4 – Client..
ShardRedundancyLevel	Shard redundancy level	The target number of redundant copies for Object data.
LocatorCount	Locator Count	The number of configured locators in the system
ServersCount	Servers Count	The number of configured servers in the system

Node Properties

vRealize Operations Manager collects properties for the vRealize Operations Manager node objects.

Table 2-41. 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 2-42. Memory Properties Collected for Node Objects

Property Key	Property Name	Description
mem|RAM	System RAM	System RAM

Table 2-43. Service Properties Collected for Node Objects

Property Key	Property Name	Description
service|proc|pid	Process ID	Process ID

Remote Collector Properties

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

Table 2-44. 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 2-45. Memory Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
mem|RAM	System RAM	System RAM

Table 2-46. Service Properties Collected for Remote Collector Objects

Property Key	Property Name	Description
service|proc|pid	Process ID	Process ID

Properties for vSAN

vRealize Operations Manager displays object properties for vSAN.

Properties for vSAN Disk Groups

The vRealize Operations Manager displays the following property for vSAN disk groups:

- vSAN Disk Groups:Configuration|vSAN Configuration

Properties for vSAN Cluster

The vRealize Operations Manager displays the following properties for vSAN cluster.

- Cluster Configuration|vSAN|Deduplication and Compression Enabled
- Cluster Configuration|vSAN|Preferred fault domain
- Cluster Configuration|vSAN|Stretched Cluster
- Cluster Configuration|vSAN|vSAN Configuration

Properties for vSAN Enabled Host

The vRealize Operations Manager displays the following property for vSAN enabled host.

- Configuration|vSAN Enabled

Properties for vSAN Cache Disk

The vRealize Operations Manager displays the following properties for vSAN cache disk.

Properties for vSAN include:

Component	Metrics
Configuration	■ Configuration Properties|Name ■ Configuration Properties|Size ■ Configuration Properties|Vendor ■ Configuration Properties|Type ■ Configuration Properties|Queue Depth
SCSI SMART Statistics	■ SCSI SMART Statistics|Media Wearout Indicator Threshold ■ SCSI SMART Statistics|Write Error Count Threshold ■ SCSI SMART Statistics|Read Error Count Threshold ■ SCSI SMART Statistics|Reallocated Sector Count Threshold ■ SCSI SMART Statistics|Raw Read Error Rate Threshold ■ SCSI SMART Statistics|Drive Temperature Threshold ■ SCSI SMART Statistics|Drive Rated Max Temperature Threshold ■ SCSI SMART Statistics|Write Sectors TOT Count Threshold ■ SCSI SMART Statistics|Read Sectors TOT Count Threshold ■ SCSI SMART Statistics|Initial Bad Block Count Threshold

Properties for vSAN Capacity Disk

The vRealize Operations Manager displays the following properties for vSAN capacity disk.

Properties for vSAN include:

Component	Metrics
Configuration	■ Configuration Properties|Name ■ Configuration Properties|Size ■ Configuration Properties|Vendor ■ Configuration Properties|Type ■ Configuration Properties|Queue Depth
SCSI SMART Statistics	■ SCSI SMART Statistics|Media Wearout Indicator Threshold ■ SCSI SMART Statistics|Write Error Count Threshold ■ SCSI SMART Statistics|Read Error Count Threshold ■ SCSI SMART Statistics|Reallocated Sector Count Threshold ■ SCSI SMART Statistics|Raw Read Error Rate Threshold ■ SCSI SMART Statistics|Drive Temperature Threshold ■ SCSI SMART Statistics|Drive Rated Max Temperature Threshold ■ SCSI SMART Statistics|Write Sectors TOT Count Threshold ■ SCSI SMART Statistics|Read Sectors TOT Count Threshold ■ SCSI SMART Statistics|Initial Bad Block Count Threshold

Alert Definitions in vRealize Operations Manager 3

Alert definitions are a combination of symptoms and recommendations that identify problem areas in vRealize Operations Manager and generate alerts on which you act for those areas.

Alert definitions are provided for various objects in your environment. You can also create your own alert definitions. See the *vRealize Operations Manager User Guide*.

- Datastore Alert Definitions on page 187

  The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

- Data Center Alert Definitions on page 192

  The vCenter adapter provides alert definitions that generate alerts on the Data Center objects in your environment.

- Custom Data Center Alert Definitions on page 193

  The vCenter adapter provides alert definitions that generate alerts on the Custom Data Center objects in your environment.

Cluster Compute Resource Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Cluster Compute Resource objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact	Health
Criticality	Symptom-based

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has CPU contention caused by less than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ > 0 descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ <= 50% of descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 Use the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has CPU contention caused by more than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ Cluster CPU demand at warning/immediate/critical level ■ > 50% of descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase CPU capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has CPU contention caused by overpopulation of virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU contention at warning/immediate/critical level ■ Cluster CPU workload at warning/immediate/critical level ■ = 0 descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase CPU capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has unexpected high CPU workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster CPU workload above DT ■ Cluster CPU workload at warning/immediate/critical level	1 Check the applications running on the virtual machines in the cluster to determine whether high CPU workload is an expected behavior. 2 Add more hosts to the cluster to increase CPU capacity. 3 Use vSphere vMotion to migrate some virtual machines to a different cluster if possible.
Fully-automated DRS-enabled cluster has memory contention caused by less than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ > 0 descendant virtual machines have [Virtual machine memory workload at warning /immediate/critical level] ■ <= 50% of descendant virtual machines have [Virtual machine memory workload at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Fully-automated DRS-enabled cluster has memory contention caused by more than half of the virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level ■ > 50% of descendant virtual machines have [Virtual machine memory demand at warning/ immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Fully-automated DRS-enabled cluster has memory contention caused by overpopulation of virtual machines.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level ■ = 0 descendant virtual machines have [Virtual machine memory demand at warning /immediate/critical level] ■ DRS Migration Threshold is not zero	1 Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads. 2 User the workload balance feature in vRealize Operations to migrate one or more virtual machines to a different cluster. 3 Use vMotion to migrate some virtual machines to a different cluster if possible. 4 Add more hosts to the cluster to increase memory capacity. 5 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
More than 5% of virtual machines in the cluster have memory contention due to memory compression, ballooning or swapping.	■ Virtual machine memory limit is set AND ■ > 5% of descendant virtual machines have [virtual machine memory contention is at warning/immediate/critical level] AND ■ > 5% of descendant virtual machines have [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/immediate/critical level]	1 Add more hosts to the cluster to increase memory capacity. 2 Use vMotion to migrate some virtual machines off the host or cluster.
Fully-automated DRS-enabled cluster has unexpected high memory workload and contention.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ Cluster memory contention above DT ■ Cluster memory content is at warning/immediate/critical level ■ Cluster memory workload at warning/immediate/critical level	1 Check the applications running on the virtual machines in the cluster to determine whether high memory workload is an expected behavior. 2 Add more hosts to the cluster to increase memory capacity. 3 Use vSphere vMotion to migrate some virtual machines to a different cluster if possible.

Alert Definition	Symptoms	Recommendations
vSphere HA failover resources are insufficient.	vSphere HA failover resources are insufficient (fault symptom)	To resolve this problem, use similar CPU and memory reservations for all virtual machines in the cluster. If this solution is not possible, consider using a different vSphere HA admission control policy, such as reserving a percentage of cluster resource for failover. Alternatively, you can use advanced options to specify a cap for the slot size. For more information, see the vSphere Availability Guide. Hosts that have vSphere HA agent errors are not good candidates for providing failover capacity in the cluster and their resources are not considered for vSphere HA admission control purposes. If many hosts have a vSphere HA agent error, vCenter Server generates this event leading to the fault. To resolve vSphere HA agent errors, check the event logs for the hosts to determine the cause of the errors. After you resolve any configuration problems, reconfigure vSphere HA on the affected hosts or on the cluster.
vSphere HA master missing.	vCenter Server is unable to find a master vSphere HA agent (fault symptom)	Check the fault page under the **Analysis** tab for this object to find more objects.

Host System Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Host System objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Symptom-based

Alert Definition	Symptoms	Recommendations
Standalone host has CPU contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] ■ <= 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level]	Use 1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has CPU contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ Host CPU demand at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has CPU contention due to overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host CPU contention is at warning/immediate/critical level ■ Host CPU demand at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Host in a cluster that does not have fully-automated DRS enabled has contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention is at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] ■ <= 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention at warning/immediate/critical level ■ Host CPU demand at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host CPU contention at warning/immediate/critical level ■ Host CPU demand at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Standalone host has memory contention caused by less than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has memory contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Standalone host has memory contention due to overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/critical level]	1 Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.

Alert Definition	Symptoms	Recommendations
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by less than half of the virtual machines.	Symptoms include the following: ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory contention at warning/immediate/critical level ■ > 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/critical level] ■ <= 50% of child virtual machines have [Virtual machine memory workload at warning/ immediate/critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by more than half of the virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by overpopulation of virtual machines.	Symptoms include the following: ■ Host inside a cluster ■ [DRS Enabled OR ! DRS fully automated] ■ Host memory workload at warning/immediate/critical level ■ Host memory contention at warning/immediate/critical level ■ = 0 child virtual machines have [Virtual machine memory workload at warning /immediate/critical level]	1 Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster. 2 Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity. 3 Upgrade the host to use a host that has larger memory capacity. 4 Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within vRealize Operations for recommended rightsizing of VMs.
Host is experiencing high number of received or transmitted packets dropped.	Symptoms include the following: ■ Host network received packets dropped ■ Host network transmitted packets dropped	1 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic. 2 Verify the health of the physical network adapter, configuration, driver and firmware versions. 3 Contact VMware support.

Alert Definition	Symptoms	Recommendations
Host is experiencing high number of received packets dropped.	Symptoms include the following: ■ Host network received packets dropped ■ Host network received packets dropped above DT ■ Host network data receive workload at Warning level ■ Host network data receive workload above DT ■ Host CPU demand at Critical level	1 If the host has one CPU, upgrade the host or use a host that has larger CPU capacity. 2 Add an additional NIC to the host. 3 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.
Host is experiencing high number of transmitted packets dropped.	Symptoms include the following: ■ Host network transmitted packets dropped ■ Host network transmitted packets dropped above DT ■ Host network data transmit workload at Warning level ■ Host network data transmit workload above DT ■ Host is dropping high percentage of packets	1 Add an additional NIC to the host. 2 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.
ESXi host has detected a link status 'flapping' on a physical NIC.	Physical NIC link state flapping (fault symptom).	ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert.
ESXi host has detected a link status down on a physical NIC.	Physical NIC link state down (fault symptom).	ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert.
Battery sensors are reporting problems.	Symptoms include the following: ■ Battery sensor health is red OR ■ Battery sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Baseboard Management Controller sensors are reporting problems.	Symptoms include the following: ■ Baseboard Management Controller sensor health is red OR ■ Baseboard Management Controller sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Fan sensors are reporting problems.	■ Fan sensor health is red OR ■ Fan sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.

Alert Definition	Symptoms	Recommendations
Hardware sensors are reporting problems.	■ Hardware sensor health is red OR ■ Hardware sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Memory sensors are reporting problems.	■ Memory sensor health is red OR ■ Memory sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Power sensors are reporting problems.	■ Power sensor health is red OR ■ Power sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Processor sensors are reporting problems.	■ Processor sensor health is red ■ Processor sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
SEL sensors are reporting problems.	■ SEL sensor health is red OR ■ SEL sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Storage sensors are reporting problems.	■ Storage sensor health is red OR ■ Storage sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
System Board sensors are reporting problems.	■ System board sensor health is red OR ■ System board sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.

Alert Definition	Symptoms	Recommendations
Temperature sensors are reporting problems.	■ Temperature sensor health is red OR ■ Temperature sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.
Voltage sensors are reporting problems.	■ Voltage sensor health is red OR ■ Voltage sensor health is yellow	Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Critical

Alert Definition	Symptoms	Recommendations
Host has lost connection to vCenter.	Host disconnected from vCenter	Click "Open Host in vSphere Web Client" in the Actions menu at the top of Alert details page to connect to the vCenter managing this host and manually reconnect the host to vCenter Server. After the connection to the host is restored by vCenter Server, the alert will be canceled.
vSphere High Availability (HA) has detected a network-isolated host.	vSphere HA detected a network isolated host (fault symptom).	Resolve the networking problem that prevents the host from pinging its isolation addresses and communicating with other hosts. Make sure that the management networks that vSphere HA uses include redundancy. With redundancy, vSphere HA can communicate over more than one path, which reduces the chance of a host becoming isolated.
vSphere High Availability (HA) has detected a possible host failure.	vSphere HA detected a host failure (fault symptom).	Find the computer that has the duplicate IP address and reconfigure it to have a different IP address. This fault is cleared and the alert canceled when the underlying problem is resolved, and the vSphere HA master agent is able to connect to the HA agent on the host. **Note** You can use the Duplicate IP warning in the `/var/log/vmkernel` log file on an ESX host or the /var/log/messages log file on an ESXi host to identify the computer that has the duplicate IP address.

Alert Definition	Symptoms	Recommendations
Host is experiencing network contention caused by too much traffic.	Symptoms include all the following: ■ Host is experiencing dropped network packets ■ Host network workload at warning/immediate/critical level	1 Review the load balancing policy in the Port Group and the vSwitch. 2 Add an additional NIC to the host. 3 Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.
The host has lost connectivity to a dvPort.	Lost network connectivity to dvPorts (fault symptom).	Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the dvPort.

Alert Definition	Symptoms	Recommendations
The host has lost connectivity to the physical network.	Lost network connectivity (fault symptom).	To determine the actual failure or to eliminate possible problems, check the status of the vmnic in the vSphere Client or from the ESX service console: ■ To check the status in the vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently down. ■ From the service console, run the command:`esxcfg-nics`. The output that appears is similar to the following: Name PCI Driver Link Speed Duplex Description ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, bring

Alert Definition	Symptoms	Recommendations
		down each ESX host port on the physical switch, run `esxcfg-nics -l`", and observe the affected vmnics. Verify that the vmnic identified in the alert is still connected to the switch and configured properly: ■ Make sure that the network cable is still connected to the switch and to the host. ■ Make sure that the switch is connected to the system, is still functioning properly, and has not been inadvertently misconfigured. For more information, see the switch documentation. ■ Check for activity between the physical switch and the vmnic. You can check activity by performing a network trace or observing activity LEDs. ■ Check for network port settings on the physical switch. To reconfigure the service console IP address if the affected vmnic is associated with a service console, see http://kb.vmware.com/kb/1000258 If the problem is caused by your hardware, contact your hardware vendor for replacement hardware.
The host lost connectivity to a Network File System (NFS) server.	Lost connection to NFS server (fault symptom).	1 Verify the NFS server is running. 2 Check the network connection to make sure the ESX host can connect to the NFS server. 3 Determine whether the other hosts that use the same NFS mount are experiencing the same problem, and check the NFS server status and share points. 4 Make sure that you can reach the NFS server by logging into the service console and using `vmkping` to ping the NFS server: "vmkping <nfs server>". 5 For advanced troubleshooting information, seehttp://kb.vmware.com/kb/1003967
A fatal error occurred on a PCIe bus during system reboot.	A fatal PCIe error occurred.	Check and replace the PCIe device identified in the alert as the cause of the problem. Contact the vendor for assistance.
A fatal memory error was detected at system boot time.	A fatal memory error occurred.	Replace the faulty memory or contact the vendor.

Health/Immediate

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Immediate

Alert Definition	Symptom	Recommendations
The host has lost redundant connectivity to a dvPort.	Lost network redundancy to DVPorts (fault symptom).	Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the DVPort.
The host has lost redundant uplinks to the network.	Lost network redundancy (fault symptom).	To determine the actual failure or to eliminate possible problems, first connect to ESX through SSH or the console: 1 Identify the available uplinks by running `esxcfg-nics -l`. 2 Remove the reported vmnic from the port groups by running `esxcfg-vswitch -U <affected vmnic#>` affected vSwitch. 3 Link available uplinks to the affected port groups by running `esxcfg-vswitch -L <available vmnic#>` affected vSwitch. Next, check the status of the vmnic in vSphere Client or the ESX service console: 1 In vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently unavailable. 2 From the service console, run `esxcfg-nics -l`. The output that appears is similar to the following example: Name PCI Driver Link Speed Duplex Description. ---------- ---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, shut down each ESX host port on the physical switch, run the "esxcfg-nics -l" command, and observe the affected vmnics. Verify that the vmnic identified in the alert is still connected to the switch and configured properly: 1 Make sure that the network cable is still connected to the switch and to the host.

Alert Definition	Symptom	Recommendations
		2 Make sure that the switch is connected to the system, is still functioning properly, and was not inadvertently misconfigured. (See the switch documentation.) 3 Perform a network trace or observe activity LEDs to check for activity between the physical switch and the vmnic. 4 Check for network port settings on the physical switch. If the problem is caused by hardware, contact your hardware vendor for a hardware replacement.
A PCIe error occurred during system boot, but the error is recoverable.	A recoverable PCIe error occurred.	The PCIe error is recoverable, but the system behavior is dependent on how the error is handled by the OEM vendor's firmware. Contact the vendor for assistance.
A recoverable memory error has occurred on the host.	A recoverable memory error occurred.	Since recoverable memory errors are vendor-specific, contact the vendor for assistance.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
ESXi Host is violating vSphere 5.5 Hardening Guide.	■ Active directory authentication disabled OR ■ Non-compliant NTP service startup policy OR ■ SSH service is running OR ■ NTP service stopped OR ■ Non-compliant timeout value for automatically disabling local and remote shell access OR ■ vSphere Authentication Proxy not used for password protection when adding ESXi hosts to active directory OR ■ Persistent logging disabled OR ■ Bidirectional CHAP for iSCSI traffic disabled OR ■ Non-compliant firewall setting to restrict access to NTP client OR ■ NTP server for time synchronization not configured OR ■ Non-compliant ESXi Shell service startup policy OR ■ Non-compliant firewall setting to restrict access to SNMP server OR ■ ESXi Shell service is running OR ■ Non-compliant DCUI service startup policy OR ■ Dvfilter bind IP address configured OR ■ Non-compliant SSH service startup policy OR ■ DCUI service is running OR ■ Non-compliant idle time before an interactive shell is automatically logged out OR ■ Non-compliant DCUI access user list OR ■ Remote syslog is not enabled	Fix the vSphere 5.5 Hardening Guide Rules Violations according to the recommendations in the vSphere5 Hardening Guide

vRealize Automation Alert Definitions

Alert definitions are combinations of symptoms and recommendations that identify problem areas in your environment and generate alerts on which you can act.

Symptom and alert definitions are defined for vRealize Automation objects. The alerts are population-based alerts based on the risk or health of a certain percentage of child objects. There are no alerts generated for network profiles.

The health and risk thresholds are as follows:

Health

- When 25%-50% of the child objects have health issues, the parent object will trigger an alert with a Warning health level.

- When 50%-75% of the child objects have health issues, the parent object will trigger an alert with an Immediate health level.
- When 75%-100% of the child objects have health issues, the parent object will trigger an alert with a Critical health level.

Risk

- When 25%-50% of the child objects have risk issues, the parent object will trigger an alert with a Warning risk level.
- When 50%-75% of the child objects have risk issues, the parent object will trigger an alert with an Immediate risk level.
- When 75%-100% of the child objects have risk issues, the parent object will trigger an alert with a Critical risk level.

vSAN Alerts Definitions

vRealize Operations Manager generates an alert if a problem occurs with the components in the storage area network that the vSAN adapter is monitoring.

Table 3-1. vSAN Alert Definitions

Alert	Alert Level	Affected Object Type	Description
Performance Service on vSAN cluster might be off or experience issues	Critical	vSAN Adapter Instance	Triggered when the vSphere Virtual SAN Performance Service is off or experiences issues for one of the vSAN-enabled cluster compute resources. Cleared by enabling Virtual SAN performance service in vSphere.
vSAN adapter instance failed to collect data from Virtual SAN Health Service. The health service might have issues.	Critical	vSAN Adapter Instance	Triggered when the vSAN adapter instance is unable to collect data from the vSphere Virtual SAN Health Service for one of the vSAN-enabled cluster compute resources. Check Virtual SAN health service settings in vSphere.
vSAN Cluster disk space usage is approaching capacity	Warning	Cluster Compute Resource	Triggered when the disk usage in a vSAN cluster reaches 80% of capacity. Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster.
vSAN Cluster disk space capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the disk usage in a vSAN cluster reaches 95% of capacity. Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster.
vSAN Cluster flash read cache is approaching capacity	Warning	Cluster Compute Resource	Triggered when the Read Cache (RC) in the vSAN cluster reaches 80% of capacity. Cleared by adding flash storage to the read cache.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
vSAN Cluster flash read cache capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the Read Cache (RC) in the vSAN cluster reaches 95% of capacity. Cleared by adding flash storage to the read cache.
vSAN Cluster virtual disk count is approaching capacity	Warning	Cluster Compute Resource	Triggered when the number of virtual disks per host in the vSAN cluster reaches 75% of capacity. Cleared by adding most hosts to the cluster.
vSAN Cluster virtual disk count capacity is less than 5%	Critical	Cluster Compute Resource	Triggered when the number of virtual disks per host in the vSAN cluster reaches 95% of capacity. Cleared by adding most hosts to the cluster.
vSAN Disk Group read cache hit rate is less than 90%	Warning	vSAN Disk Group	Triggered when the vSAN disk group read cache hit rate is less than 90%. Cleared by adding more cache to accommodate the workload.
vSAN Disk Group read cache hit rate is less than 90% and write buffer free space is less than 10%	Warning	vSAN Disk Group	Triggered when the vSAN disk group read cache hit rate is less than 90% and the vSAN disk group write buffer free space is less than 10%. Cleared by adding more flash capacity to the vSAN disk group.
vSAN Host has no VMkernel NIC configured	Immediate	Host System	Triggered when vSAN host has no VMkernel NIC configured. Clears when the symptom disappears.
One or more physical disks on vSAN host is experiencing software state health issues	Critical	Host System	Triggered when one or more physical disks on vSAN host is experiencing software state health issues.
vSAN Performance Service is unable to communicate and retrieve statistics from host	Critical	Host System	Triggered when vSAN Performance Service is unable to communicate and retrieve statistics from host.
vSAN enabled hosts have inconsistent values for advanced configuration options	Critical	vSAN Cluster	Triggered when some advanced configuration settings have different values on different hosts in the vSAN cluster.
vSAN is disabled on the host	Critical	Host System	Triggered when vSAN is disabled on the host. vSAN objects health will be impacted If this host stores any vSAN data (For example: virtual machine objects) on its local disks.
One or more vSAN enabled hosts are not in the same IP subnet	Critical	Host System	Triggered when one or more vSAN enabled hosts are not in the same IP subnet.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
Host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic	Critical	Host System	Triggered when host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic. NOTE Even if an ESXi host is part of the vSAN cluster, but is not contributing storage, it must still have a VMkernel NIC configured for vSAN traffic.
One or more hosts in the vSAN cluster have misconfigured multicast addresses	Critical	Host System	Triggered when one or more hosts in the vSAN cluster have misconfigured multicast addresses.
vSAN health service is not installed on the host	Critical	Host System	Triggered when vSAN health service is not installed on the host.
Host in a vSAN cluster has IP multicast connectivity issue	Critical	Host System	Triggered when host in a vSAN cluster has IP multicast connectivity issue. It means that multicast is most likely the root cause of a vSAN network partition.
Host in a vSAN cluster has connectivity issues and vCenter Server does not know its state	Critical	Host System	Triggered when host in a vSAN cluster has connectivity issues and vCenter Server does not know its state.
vSAN disk group has incorrect deduplication and compression configuration	Critical	Host System	Triggered when vSAN disk group has incorrect deduplication and compression configuration.
vSAN witness host has an invalid preferred fault domain	Critical	Host System	Triggered when vSAN witness host has an invalid preferred fault domain.
vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster	Critical	Host System	Triggered when vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster.
Host has invalid unicast agent and impacting the health of vSAN Stretched Cluster	Critical	Host System	Triggered when host has invalid unicast agent and impacting the health of vSAN Stretched Cluster. An invalid unicast agent on the host can cause a communication malfunction with the witness host.
Storage I/O controller driver is not VMware certified	Critical	Host System	Triggered when stability and integrity of vSAN may be at risk as the storage I/O controller driver is not VMware certified.
Storage I/O controller is not compatible with the VMware Compatibility Guide	Critical	Host System	Triggered when vSAN environment may be at risk as the Storage I/O controller on the ESXi hosts that are participating in a vSAN cluster are not compatible with the VMware Compatibility Guide.
vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster	Critical	Host System	Triggered when vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
Unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster	Critical	Host System	Triggered when unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster.
The preferred fault domain is not set for the witness host in a vSAN Stretched cluster	Critical	Host System	Triggered when the preferred fault domain is not set for the witness host in a vSAN Stretched cluster and affecting the operations of vSAN Stretched cluster.
vSAN Stretched cluster contains a witness host without a valid disk group	Critical	Host System	Triggered when vSAN Stretched cluster contains a witness host without a valid disk group. If the witness host does not have any disk claimed by vSAN then its fault domain is not available.
vSAN Stretched cluster has inconsistent configuration for Unicast agent	Critical	Host System	Triggered when vSAN Stretched cluster contains multiple unicast agents. This means multiple unicast agents were set on non-witness hosts.
vSAN Stretched cluster does not contain a valid witness host	Critical	Host System	Triggered when vSAN Stretched cluster does not contain a valid witness host. This affects the operations of vSAN Stretched cluster.
vSAN cluster has multiple network partitions	Critical	Host System	Triggered when vSAN cluster has multiple network partitions due to a network issue.
Witness host is a part of vSAN Stretched cluster	Critical	Host System	Triggered when witness host is a part of the vCenter cluster, which forms vSAN Stretched cluster.
Witness host resides in one of the data fault domains	Critical	Host System	Triggered when witness host resides in one of the data fault domains. This affects the operations of vSAN Stretched cluster.
vSAN cluster has unexpected hosts	Critical	Host System	Triggered when vSAN cluster has unexpected hosts.
vSAN is unable to retrieve the physical disk information from host	Critical	Host System	Triggered when vSAN is unable to retrieve the physical disk information from host. vSAN Health Service may not be working properly on this host.
vCenter Server has lost connection to a host that is part of a vSAN cluster.	Critical	Host System	Triggered when host that is part of a vSAN cluster is in disconnected state or not responding and vCenter Server does not know its state.
vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk	Critical	Host System	Triggered when vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks.	Critical	Host System	Triggered when vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.	Critical	Host System	Triggered when vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
vSAN is using a physical disk which has high congestion value	Critical	Host System	Triggered when vSAN is using a physical disk which has high congestion value. This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive.
Disk format version of one or more vSAN disks is out of date	Critical	Host System	Triggered when disk format version of one or more vSAN disks is out of date and is not compatible with other vSAN disks. This can lead to problems in creating or powering on VMs, performance degradation, and EMM failures.
vSAN Cluster has multiple Stats DB objects which are creating conflicts and affecting vSAN Performance Service	Critical	vSAN Cluster	Triggered when vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service.
vSAN cluster has issues in electing stats master of vSAN Performance service	Critical	vSAN Cluster	Triggered when vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service.
CLOMD process on the host has issues and impacting the functionality of vSAN cluster	Critical	Host System	Triggered when CLOMD process on the host has issues and impacting the functionality of vSAN cluster.
Number of vSAN components on a disk is reaching or has reached its limit	Critical	vSAN Cluster	Triggered when number of vSAN components on a disk is reaching or has reached its limit. This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
vSAN HCL DB auto updater is not working properly	Critical	vSAN Cluster	Triggered when vSAN HCL DB auto updater is not working properly. This means that vSAN cannot download and update its HCL DB automatically.
vSAN has encountered an issue while reading the metadata of a physical disk	Critical	Host System	Triggered when vSAN has encountered an issue while reading the metadata of a physical disk and cannot use this disk.
Overall health of the physical disks in a vSAN Cluster is impacted	Critical	Host System	Triggered when overall health of the physical disks in a vSAN Cluster is impacted. See the health status of each physical disk individually on all the hosts.
vSAN performance service statistics database object is reporting issues.	Critical	Host System	Triggered when vSAN performance service statistics database object is reporting issues.
Overall health of vSAN objects is reporting issues	Critical	vSAN Cluster	Triggered when overall health of vSAN objects is reporting issues.
Number of vSAN components on a host is reaching or has reached its limit	Critical	Host System	Triggered when number of vSAN components on a host is reaching or has reached its limit. This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations.
Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster	Critical	vSAN Cluster	Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster.
Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster	Critical	Host System	Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster.
Statistics collection of vSAN performance service is not working correctly	Critical	vSAN Cluster	Triggered when statistics collection of vSAN performance service is not working correctly. This means that statistics collection or writing statistics data to storage have failed for three consecutive intervals.
After one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects	Critical	vSAN Cluster	Triggered when after one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects.
vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations	Critical	vSAN Cluster	Triggered when vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations.

Table 3-1. vSAN Alert Definitions (Continued)

Alert	Alert Level	Affected Object Type	Description
Disk load variance between some vSAN disks exceeded the threshold value	Critical	vSAN Cluster	Triggered when disk load variance between some vSAN disks exceeded the threshold value. vSAN cannot perform the load balance properly.
Host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host	Critical	Host System	Triggered when host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host.
Storage I/O controller drivers is not supported with the current version of ESXi running on the host	Critical	Host System	Triggered when stability and integrity of vSAN may be at risk as the storage I/O controller driver is not supported with the current version of ESXi running on the host.
vSAN HCL DB is not up-to-date	Critical	vSAN Cluster	Triggered when vSAN HCL DB is not up-to-date.
vSAN cluster health checks are reporting issues	Critical	vSAN Cluster	Triggered when vSAN cluster health checks are reporting issues.
vSAN cluster Hardware Compatibility health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster Hardware Compatibility health checks are experiencing issues.
vSAN cluster Limits health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster Limits health checks are experiencing issues.
vSAN cluster network health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN cluster network health checks are experiencing issues.
Performance Service on vSAN cluster might be off or experiencing issues	Critical	vSAN Cluster	Triggered when performance service on vSAN cluster might be off or experiencing issues. vROps will not be able to collect the data from vSAN cluster in this state.
vSAN Stretched cluster health checks are experiencing issues	Critical	vSAN Cluster	Triggered when vSAN Stretched cluster health checks are experiencing issues.
MTU check (ping with large packet size) has failed on vSAN host	Critical	Host System	Triggered when MTU check (ping with large packet size) has failed on vSAN environment due to some MTU misconfiguration in the vSAN network.
Basic (unicast) connectivity check (normal ping) has failed on vSAN host	Critical	Host System	Triggered when basic (unicast) connectivity check (normal ping) has failed on vSAN host due to network misconfiguration.
vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.	Critical	Host System	Triggered when vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues.

Alerts in the vSphere Web Client

The vSphere Web Client displays the results of health tests for the following vSAN monitored groups:

- Network
- Physical disk
- Cluster
- Limits
- Data
- Hardware compatibility
- Performance Service
- Stretched Cluster (if enabled)

Each group contains several individual checks. If a check fails, the vSAN adapter issues a warning or error level alert. The alert indicates the host or cluster where the problem occurred and provides a recommendation to clear the alert. For a complete list of all vSAN health test alerts, see Knowledge Base article 2114803.

vSphere Distributed Port Group

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Port objects in your environment.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Critical

Alert Definition	Symptom	Recommendations
One or more ports are in link down state.	Symptoms include all of the following: ■ Port is connected ■ One or more ports are in a link down state	Verify that there is physical connectivity for the NICs on the host. Verify the admin status on the port.
One or more ports are experiencing network contention.	Port is experiencing dropped packets	Check if the packet drops are due to high CPU resource utilization or uplink bandwidth utilization. User vMotion to migrate the virtual machine that the port is attached to a different host.

Virtual Machine Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the virtual machine objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine is experiencing memory compression, ballooning or swapping due to memory limit.	■ Virtual machine memory limit is set AND ■ Virtual machine memory demand exceeds configured memory limit AND ■ [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/immediate/critical level] AND ■ Recommended virtual machine memory size	Increase the memory limit for the virtual machine to match the recommended memory size. Alternatively, remove memory limit for the virtual machine.
Virtual machine has CPU contention caused by swap wait.	Virtual machine CPU swap wait is at warning/Immediate/Critical level.	1 Upgrade the host with more memory. 2 Use vSphere vMotion to migrate this virtual machine to a different host or cluster. 3 Set memory reservations for the virtual machine to prevent swapping.
Virtual machine has CPU contention caused by IO wait.	Virtual machine CPU I/O wait is at warning/immediate/critical level.	Increase the datastore I/O capacity for the connected data stores to reduce CPU I/O wait on the virtual machine.
Virtual machine has unexpected high CPU workload.	Symptoms include all of the following: ■ Virtual machine CPU demand at warning/immediate/critical level ■ Anomaly is starting to/moderately/critically high	1 Check the guest applications to determine whether high CPU workload is an expected behavior. 2 Add more CPU capacity for this virtual machine.
Virtual machine has unexpected high memory workload.	Symptoms include all of the following: ■ Virtual machine memory workload is at Warning/Immediate/Critical level ■ Anomaly is starting to/moderately/critically high	1 Check the guest applications to determine whether high memory workload is an expected behavior. 2 Add more memory for this virtual machine.

Alert Definition	Symptom	Recommendations
Virtual machine has memory contention due to swap wait and high disk read latency.	Symptoms include all of the following: ■ Virtual machine CPU swap wait is at warning/immediate/critical level (5/10/15) ■ Virtual machine has read latency at warning level ■ Recommended virtual machine memory size	Add more memory for this virtual machine.
Virtual machine has memory contention due to memory compression, ballooning or swapping.	■ ! Virtual machine memory limit is set AND ■ Virtual machine has memory contention at warning/immediate/critical level AN ■ [Virtual machine memory ballooning at warning/immediate/critical level OR ■ Virtual machine memory is compressed OR ■ Virtual machine is using swap]	1 Add memory reservations to this virtual machine to prevent ballooning and swapping. 2 Use vSphere vMotion to migrate this virtual machine to a different host or cluster.
Virtual machine has unexpected high disk I/O workload.	Symptoms include all of the following: ■ Virtual machine disk I/O workload at Warning/Immediate/Critical level (80/90/95) ■ Virtual machine disk I/O workload above DT	1 Check the applications running on the virtual machine to determine whether high disk I/O workload is an expected behavior. 2 Use vSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS.
Virtual machine has disk I/O read latency problem.	Symptoms include all of the following: ■ Virtual machine disk read latency at Warning /Immediate/Critical level ■ Virtual machine disk read latency above DT ■ Virtual machine has low co-stop ■ Virtual machine has low CPU swap wait	1 Check whether you have enabled Storage IO control on the datastores connected to the virtual machine. 2 Increase IOPS for the datastores connected to the virtual machine. 3 UsevSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS.
Virtual machine has disk I/O write latency problem.	Symptoms include all of the following: ■ Virtual machine disk write latency at Warning/ Immediate/Critical level ■ Virtual machine disk write latency above DT ■ Virtual machine has low CPU swap wait (< 3 ms)	1 Check whether you have enabled Storage IO Control on the data stores connected to the datastore. 2 Increase IOPS for the data stores connected to the virtual machine. 3 If the virtual machine has multiple snapshots, delete the older snapshots. 4 Use vSphere Storage vMotion to migrate some virtual machines to a different datastore.

Alert Definition	Symptom	Recommendations
Virtual machine has disk I/O latency problem caused by snapshots.	Symptoms include all of the following: ■ Virtual machine CPU I/O wait is at warning/immediate/critical level ■ Virtual machine has at least one snapshot ■ All child datastores have [! Disk command latency at warning level]	1 If the virtual machine has multiple snapshots, delete the older snapshots. 2 Reduce the number of snapshots by consolidating the snapshots into one snapshot. In vSphere Client, select the VM, right-click, select **Snapshot**, and then **Consolidate**.
Virtual machine is consuming disk space in a rapid and unexpected manner.	Symptoms include all of the following: ■ Guest file system overall disk space usage reaching warning/immediate/critical limit (80, 90, 95) ■ Virtual machine disk space time remaining high (> 60 days) ■ Guest file system space usage above DT ■ Guest partition disk space usage	1 Check the application and verify that it is behaving correctly. 2 Add a new hard disk to the virtual machine and configure the guest file system partition to use the disk.
One or more guest file systems is out of disk space.	One or more guest file systems out of disk space (Fault symptom).	Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk.
Not enough resources for vSphere HA to start the virtual machine.	Not enough resources forvSphere HA to start VM (Fault symptom).	1 If virtual machine CPU reservation is set, decrease the CPU reservation configuration. 2 If virtual machine memory reservation is set, decrease the memory reservation configuration. 3 Add more hosts to cluster. 4 Bring any failed hosts online or resolve a network partition, if one exists. 5 If DRS is in manual mode, look for pending recommendations and approve the recommendations so that vSphere HA failover can proceed.
The Fault tolerance state of the virtual machine has changed to "Disabled" state.	VM fault tolerance state changed to disabled (Fault symptom).	Enable the secondary virtual machine indicated in the alert.
vSphere HA failed to restart a network isolated virtual machine.	vSphere HA failed to restart a network isolated virtual machine (Fault symptom).	Manually power on the virtual machine.
The fault tolerance state of the virtual machine has changed to "Needs Secondary" state.	VM Fault Tolerance state changed to needs secondary (Fault symptom).	Keep HA enabled when Fault tolerance (FT) is required to protect virtual machines.

Alert Definition	Symptom	Recommendations
vSphere HA cannot perform a failover operation for the virtual machine	vSphere HA virtual machine failover unsuccessful (Fault symptom)	1 If the error information reports that a file is locked, the virtual machine might be powered on a host that the vSphere HAmaster agent can no longer monitor by using the management network or heartbeat datastores. 2 The virtual machine might have been powered on by a user on a host outside of the cluster. If any hosts are declared offline, determine whether a networking or storage problem caused the situation. 3 If the error information reports that the virtual machine is in an invalid state, an in-progress operation might be preventing access to the virtual machine files. Determine whether any operations are in progress, such as a clone operation that is taking a long time to complete. 4 You can also try to power on the virtual machine and investigate any returned errors.
Virtual machine is experiencing memory compression, ballooning or swapping due to memory limit.	■ Virtual machine memory limit is set ■ Virtual machine memory demand exceeds configured memory limit ■ [Virtual machine memory is compressed OR ■ Virtual machine is using swap OR ■ Virtual machine memory ballooning is at warning/immediate/critical level] ■ Recommended virtual machine memory size	Increase the memory limit for the virtual machine to match the recommended memory size. Alternatively, remove memory limit for the virtual machine.

Efficiency/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Efficiency

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine has large disk snapshots.	Symptoms include all of the following: ■ Virtual machine has large disk snapshots ■ Reclaimable snapshot waste ■ Datastore space usage reaching warning/immediate/critical limit	If the virtual machine has multiple snapshots, delete the older snapshots.

Efficiency/Warning

These alert definitions have the following impact and criticality information.

Impact Efficiency

Criticality Warning

Alert Definition	Symptom	Recommendations
Virtual machine is idle.	Symptoms include all of the following: ■ Virtual machine is idle ■ Virtual machine high ready time on each vCPU ■ ! Virtual machine is powered off	Power off this virtual machine to allow for other virtual machines to use CPU and memory that this virtual machine is wasting.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
Virtual machine has CPU contention caused by co-stop.	Symptoms include all of the following: ■ Virtual machine CPU co-stop at warning/immediate/critical level ■ ! Virtual machine is powered off ■ Number of vCPUs to remove from virtual machine	Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the symptom.
Virtual machine has chronic high CPU workload leading to CPU stress.	Symptoms include all of the following: ■ Virtual machine CPU stress is at warning/immediate/critical level ■ Recommended number of vCPUs to add	Add more CPU capacity for this virtual machine.
Virtual machine has high CPU co-stop due to snapshots.	Symptoms include all of the following: ■ Virtual machine CPU co-stop is at warning/immediate/critical level ■ Virtual machine has at least one snapshot	To reduce the high co-stop (%CSTP) values and increase virtual machine performance, consolidate any snapshots into the main virtual disk. In the vSphere Client, select the VM, right click, and select **Snapshot**, and then **Consolidate**. After consolidation, the %CSTP value is reduced or eliminated and VM performance is improved. If performance is not improved enough, continue researching other potential VM performance issues. See VMware KB: http://kb.vmware.com/kb/2000058

Alert Definition	Symptom	Recommendations
Virtual machine has chronic high memory workload leading to memory stress.	Symptoms include all of the following: ■ Virtual machine memory stress at warning/immediate/critical level ■ Recommended virtual machine memory size > 0	Add more memory for the VM.
Virtual machine is projected to run out of disk space.	Symptoms include all of the following: ■ Virtual machine disk space time remaining low (<= 60 days) ■ ! Guest file system space usage above DT ■ ! Guest file system overall disk space usage reaching warning limit (85%) ■ Guest partition disk space usage	1 Check the application configuration to determine whether the virtual machine disk capacity will be sufficient. 2 Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk.

Alert Definition	Symptom	Recommendations
Virtual machine is running out of disk space.	Symptoms include all of the following: ■ Guest file system overall disk space usage reaching warning/immediate/critical limit (80, 90, 95) ■ Virtual machine disk space time remaining low (<= 60 days) ■ ! Guest file system space usage above DT ■ Guest partition disk space usage	1 Add a new hard disk to the virtual machine and configured the guest file system partition to use the disk. 2 Reclaim disk space using in-guest disk cleanup mechanisms.
Virtual machine is violating vSphere 5.5 hardening guide.	■ Unrestricted VM-to-VM communication through VMCI OR ■ VMsafe CPU/Memory APIs-port number configured OR ■ Dvfilter network API enabled OR ■ Non-compliant max VMX file size OR ■ Non-compliant max VM log file size OR ■ Allow unauthorized modification of device settings OR ■ Allow unauthorized connect and disconnect of devices OR ■ Tools auto install not disabled OR ■ Non-compliant max number of remote console connections OR ■ Allow VM to obtain detailed information about the physical host OR ■ Non-compliant max VM log file count OR ■ Feature not exposed in vSphere: MemsFss is not disabled OR ■ VMsafe CPU/memory API enabled OR ■ Parallel port connected OR ■ Console drag and drop operation not disabled OR ■ Console copy operation not disabled OR ■ Serial port connected OR ■ Feature not exposed in vSPhere: AutoLogon is not disabled OR ■ Use independent non persistent disk OR ■ Feature not exposed in vSphere: UnityPush is not disabled OR ■ Shrink virtual disk not disabled - diskShrink OR ■ Feature not exposed in vSphere: GetCreds is not disabled OR ■ CD-ROM connected OR ■ Feature not exposed in vSPhere: HGFSServerSet is not disabled OR	Fix the vSphere 5.5 hardening guide rule violations according to the recommendations in the vSphere Hardening Guide (XLSX).

Alert Definition	Symptom	Recommendations
	■ 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 ■ Feature not exposed in vSphere: Unity-Unityactive is not disabled	

Risk/Warning

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Warning

Alert Definition	Symptom	Recommendations
Virtual machine is demanding more CPU than the configured limit.	Symptoms include all of the following: ■ Virtual machine CPU limit is set ■ Virtual machine CPU demand exceeds configured limit ■ ! Virtual machine's CPU demand exceeds its provisioned capacity	Increase of remove CPU limits on the VM.

vSphere Distributed Switch Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Switch objects in your environment.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Critical

Alert Definition	Symptom	Recommendations
Network traffic is blocked for one or more ports.	Network traffic is blocked for one or more ports.	Check the security policy on the port groups as well as any ACL rule configuration.

Health/Warning

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Warning

Alert Definition	Symptom	Recommendations
Distributed Switch configuration is out of sync.	Distributed Switch configuration is out of sync with the vCenter Server.	Change the distributed switch configuration to match the host. Identify the distributed switch properties that are out of sync. If these properties were changed locally on the host in order to maintain connectivity, update the distributed switch configuration in the vCenter Server. Otherwise, re-apply the the vCenter Server configuration to this host.
One or more VLANs are unsupported by the physical switch.	One or more VLANs are unsupported by the physical switch.	Ensure the VLAN configuration on the physical switch and the distributed port groups are consistent.
Teaming configuration does not match the physical switch.	Teaming configuration does not match the physical switch.	Ensure the teaming configuration on the physical switch and the distributed switch are consistent.

Alert Definition	Symptom	Recommendations
The MTU on the Distributed Switch is not allowed by one or more VLANs on the host.	The MTU on the Distributed Switch is not allowed by one or more VLANs on the host.	Ensure the MTU configuration on the physical switch and the distributed switch are consistent.
There is an MTU mismatch between the host and a physical switch.	There is an MTU mismatch between the host and a physical switch.	Adjust the MTU configuration on the host to match the physical switch. Change the MTU configuration on the physical switch.

Risk/Warning

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Warning

Alert Definition	Symptom	Recommendations
The distributed switch configuration is incorrect.	Host without redundant physical connectivity to the distributed switch.	Verify that at least two NICs on each host is connected to the distributed switch.

vCenter Server Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vCenter Server objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
A problem occurred with a vCenter Server component.	The vCenter Server health changed (fault symptom).	The actions to take to resolve the problems depend on the specific problem that caused the fault. Review the issue details, and check the documentation.
Duplicate object name found in thevCenter Server.	Duplicate object name found in the vCenter Server.	Ensure the virtual machines names are unique before enabling the Name-Based Identification feature.
The vCenter Server Storage data collection failed.	The vCenter Server storage data collection failed.	Ensure vCenter Management Webservice is started and Storage Management Service is functioning.

Datastore Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

Health/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
Datastore has unexpected high Disk I/O workload.	Symptoms include all of the following: ■ Datastore disk I/O workload at warning/immediate/critical level ■ Datastore disk I/O workload above DT	1 Check the applications running on the virtual machines placed on the datastore to determine whether high disk I/O workload is expected behavior. 2 Increase IOPS for the datastore.
Datastore is consuming disk space in a rapid and unexpected manner.	Symptoms include all of the following: ■ Datastore space usage reaching warning/immediate/critical level ■ Datastore space growth above DT ■ Datastore time remaining high	1 Check if there is an unexpected provisioning of virtual machines on this datastore. 2 Use vSphere Storage vMotion to migrate some virtual machines to a different datastore. 3 Add more capacity to the datastore.

Health/Critical

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Critical

Alert Definition	Symptom	Recommendations
A storage device for a datastore has been detected to be off.	Storage device has been turned off administratively (fault symptom).	Ask the administrator about the device state. The fault will be resolved and the alert canceled if the device is turned on. If SCSI devices are detached or permanently removed, you must manually cancel the alert.
Datastore has lost connectivity to a storage device.	Host(s) lost connectivity to storage device(s) (fault symptom).	The storage device path, for example, `vmhba35:C1:T0:L7`, contains several potential failure points: Path Element | Failure Point ---------- vmhba35 | HBA (Host Bus Adapter) C1 | Channel T0 | Target (storage processor port) L7 | LUN (Logical Unit Number or Disk Unit). To determine the cause of the failure or to eliminate possible problems: Identify the available storage paths to the reported storage device by running `esxcfg-mpath -l`. For more information, see 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. Troubleshoot the connectivity to the iSCSI storage by using the software initiator: 1 Check whether a ping to the storage array fails from ESX. For more information, see http://kb/vmware.com/kb/1003486 2 Check whether a vmkping to each network portal of the storage array fails. For more information, seehttp://kb.vmware.com/kb/1003728. 3 Check that the initiator is registered on the array. For more information, contact your storage vendor. 4 Check that the following physical hardware is functioning correctly: Ethernet switch, Ethernet cables between the switch and the ESX host, and Ethernet cables between the switch and the storage array. 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

Alert Definition	Symptom	Recommendations
		messages to the ESX hosts. For more information about configuring the fiber switch, see http://kb.vmware.com/kb/1002301. Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself. You must rescan after making changes to make sure that the targets are detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or change, you must cancel the fault alert as a workaround. The alert will then be canceled automatically.

Health/Immediate

These alert definitions have the following impact and criticality information.

Impact Health

Criticality Immediate

Alert Definition	Symptom	Recommendations
Datastore has one or more hosts that have lost redundant paths to a storage device.	Host(s) lost redundancy to storage device(s) (fault symptom).	The storage device path, for example, vmhba35:C1:T0:L7, contains several potential failure points: Path Element | Failure Point ---------- vmhba35 | HBA (Host Bus Adapter) C1 | Channel T0 | Target (storage processor port) L7 | LUN (Logical Unit Number or Disk Unit). Use the following guidance to determine the cause of the failure or to eliminate possible problems. Identify the available storage paths to the reported storage device by running `esxcfg-mpath -l`. For more information, see 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. Troubleshoot the connectivity to the iSCSI storage by using the software initiator: 1 Check whether a ping to the storage array fails from ESX. For more information, seehttp://kb/vmware.com/kb/1003486. 2 Check whether a vmkping to each network portal of the storage array fails. For more information, seehttp://kb.vmware.com/kb/1003828. 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. 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, seehttp://kb.vmware.com/kb/1002301.

Alert Definition	Symptom	Recommendations
		Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself. You must rescan after making changes to make sure that the targets are detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or change, you must cancel the fault alert as a workaround. The alert will be canceled automatically after that.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Symptom-based

Alert Definition	Symptom	Recommendations
Datastore is running out of disk space.	Symptoms include all of the following: ■ Datastore space usage reaching warning/immediate/critical level ■ ! Datastore space growth above DT ■ Datastore space time remaining is low	1 Add more capacity to the datastore. 2 Use vSphere vMotion to migrate some virtual machines to a different datastore. 3 Delete unused snapshots of virtual machines from datastore. 4 Delete any unused templates on the datastore.
Datastore is projected to run out of disk space.	Symptoms include all of the following: ■ ! Datastore space usage reaching warning level ■ ! Datastore space growth above DT ■ Datastore space time remaining is low	1 Check if datastore usage is a planned growth and expand the storage if necessary. 2 Use vSphere vMotion to migrate some virtual machines to a different datastore.

Data Center Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Data Center objects in your environment.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information:

Impact Risk

Criticality Symptom-based

Alert Definition	Symptoms	Recommendations
Data center has unbalanced CPU "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ DC is unbalanced on CPU "demand" workload ■ DC has significant CPU "demand" workload difference ■ At least one cluster in DC has high CPU "demand" workload	Rebalance the container to spread the workload more evenly.
Data center has unbalanced memory "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully enabled ■ DC is unbalanced on memory "demand" workload difference ■ At least one cluster in DC has high memory "demand" workload	Rebalance the container to spread the workload more evenly.
Data center has unbalanced memory "consumed" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ DC is unbalanced on memory "consumed" workload ■ DC has significant memory "consumed" workload difference ■ At least one cluster in DC has high memory "consumed" workload	Rebalance the container to spread the workload more evenly.

Custom Data Center Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Custom Data Center objects in your environment.

Risk/Symptom-Based

These alert definitions have the following impact and criticality information.

Impact Risk

Criticality Symptom-based

Alert Definition	Symptoms	Recommendations
Custom data center has unbalanced CPU "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on CPU "demand" workload ■ CDC has significant CPU "demand" workload difference ■ At least one cluster in CDC has high CPU "demand" workload	Rebalance the container to spread the workload more evenly.
Custom data center has unbalanced memory "demand" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on memory "demand" workload ■ CDC has significant memory "demand" workload difference ■ At least one cluster in CDC has high memory"demand" workload	Rebalance the container to spread the workload more evenly.
Custom Datacenter has unbalanced memory "consumed" workload.	Symptoms include all of the following: ■ DRS enabled ■ DRS fully automated ■ CDC is unbalanced on memory "consumed" workload ■ CDC has significant memory "consumed" workload difference ■ At least one cluster in CDC has high memory"consumed" workload	Rebalance the container to spread the workload more evenly.

Index